Pixel-Level Tissue Classification for Ultrasound Images

Daniel V. Pazinato, Bernardo V. Stein, Waldir R. de Almeida, Rafael de O. Werneck, Pedro R. Mendes Júnior, Otávio A. B. Penatti, Ricardo da S. Torres, Fábio H. Menezes, and Anderson Rocha

Abstract*—*Background: **Pixel-level tissue classification for ultrasound images, commonly applied to carotid images, is usually based on defining thresholds for the isolated pixel values. Ranges of pixel values are defined for the classification of each tissue. The classification of pixels is then used to determine the carotid plaque composition and, consequently, to determine the risk of diseases (e.g., strokes) and whether or not a surgery is necessary. The use of threshold-based methods dates from the early 2000s but it is still widely used for *virtual histology*. *Methodology/Principal Findings:* We propose the use of descriptors that take into account information about a neighborhood of a pixel when classifying it. We evaluated experimentally different descriptors (statistical moments, texture-based, gradient-based, local binary patterns, etc.) on a dataset of five types of tissues: blood, lipids, muscle, fibrous, and calcium. The pipeline of the proposed classification method is based on image normalization, multiscale feature extraction, including the proposal of a new descriptor, and machine learning classification. We have also analyzed the correlation between the proposed pixel classification method in the ultrasound images and the real histology with the aid of medical specialists. *Conclusions/Significance:* The classification accuracy obtained by the proposed method with the novel descriptor in the ultrasound tissue images (around 73%) is significantly above the accuracy of the state-of-the-art threshold-based methods (around 54%). The results are validated by statistical tests. The correlation between the virtual and real histology confirms the quality of the proposed approach showing it is a robust ally for the *virtual histology* in ultrasound images.**

***Index Terms*—Carotid plaque composition, pixel-level tissue classification, ultrasound images, virtual histology (VH).**

I. Introduction

EXTRACRANIAL carotid artery disease is a preventable cause of ischemic cerebrovascular accidents (strokes). The carotid plaque is usually classified according to the degree of stenosis and most surgical trials have used a 70% or greater diameter loss as indication for surgery [1]–[3]. The techniques to measure stenosis are already well established and follow international standardization. More recently, it was recognized that not only must the degree of stenosis be evaluated, but also the carotid plaque instability, as it is an important determinant of stroke risk, since it may trigger an episode of local thrombosis or lead to distal embolization of plaque debris [4], [5]. There are a number of different imaging techniques that provide information on carotid plaque morphology as the duplex ultrasound (DU), computed tomography, or magnetic resonance imaging. These techniques allow for the study of plaque surface and contents, plaque volume, and vessel wall movement [5]–[7].

The advantages of studying the plaque with DU derive from the noninvasive nature of the technique, the low cost, and easy availability of the equipment, which renders it ideal for office-based evaluations. Also, this technique is easily learned by the health professionals and the anatomic superficial location of the extracranial carotid arteries favors their insonation [5].

The study of plaque contents is directed to the identification of the plaque constituents and their relative localization to the luminal surface [8]–[10]. Lipids reflect ultrasound poorly and produce an image which is predominantly dark in the gray-scale B-mode ultrasound screen. Fibrotic tissue, which renders the plaque more stable, produces a stronger reflection of the ultrasound waves and appears lighter on the screen. For this reason, the plaques that are more vulnerable or unstable appear darker on the screen and are named *echolucent* and are also named soft plaques or complicated plaques. In opposition, more stable plaques are denser and appear lighter on the screen and are named *echogenic*.

The vulnerability or instability of the plaque determines the chances of having a plaque accident, where there may occur a plaque rupture with distal embolization of debris or local thrombosis. It is known that the presence and size of a lipid rich core, associated or not with necrotic tissue or intraplaque hemorrhage, is a determinant of plaque instability [4]. In this way, the classification of the plaque according to this necrotic/lipid core is the goal of the actual techniques.

The evaluation of the plaque constituents may be done visually by the DU operator who classifies the plaque into six categories [5]: uniformly echolucent (more than 85% of the plaque appears dark); predominantly echolucent (with 50–85% of the plaque dark); predominantly echogenic (50–85% of the plaque appears light); uniformly echogenic (more than 85% of the plaque appears light); nonclassifiable due to heavy

Manuscript received March 13, 2014; revised September 18, 2014; accepted December 19, 2014. Date of publication December 30, 2014; date of current version December 31, 2015.

D. V. Pazinato, B. V. Stein, W. R. de Almeida, R. de O. Werneck, P. R. Mendes Júnior, R. da S. Torres, and A. Rocha are with the RECOD Lab, Institute of Computing, University of Campinas, Campinas 13.083-852, SP, Brazil (e-mail: daniel.pazinato@students.ic.unicamp.br; bernardo.stein@students.ic.unicamp.br; waldir.almeida@students.ic.unicamp.br; rafael.werneck@ic.unicamp.br; pedro.junior@ic.unicamp.br; rtorres@ic.unicamp.br; anderson.rocha@ic.unicamp.br).

O. A. B. Penatti is with the Advanced Technologies Group, Samsung Research Institute, Campinas 13.083-852, SP, Brazil and also with the RECOD Lab, Institute of Computing, University of Campinas, Campinas 13.083-852, SP, Brazil (e-mail: o.penatti@samsung.com).

F. H. Menezes is with the Faculty of Medical Sciences, University of Campinas, Campinas 13.083-852, SP, Brazil (e-mail: fmenezes@fcm.unicamp.br).

This paper has supplementary downloadable material available at http://ieeexplore.ieee.org.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JBHI.2014.2386796

calcification, which precludes the visualization of the plaque contents; and plaque with a luminal surface disruption, which would indicate the presence of a plaque ulceration. This classification is subjective and has a high variability among DU operators.

In order to make it more objective, it was proposed to classify the plaques according to image and mathematical analysis of the plaque constituents. The most frequent computer-assisted analysis in ultrasound images is the *gray-scale median* (GSM), which is the median of the histogram of the pixel distribution of the plaque image according to the brightness of the pixel on a 0–255 scale. A GSM value below 32 is characteristic of plaques rich in lipids, and a GSM value below 15 is associated with a greater chance of having symptoms in patients with a carotid artery plaque [11], [12].

The precise determination of the plaque constituents and their localization is also desirable and could help in plaque classification. Lal *et al.* [9] proposed the analysis of the plaque constituents according to the pixel brightness level for *virtual histology (VH)*. The *threshold-based method* of Lal *et al.* [9] has been extensively applied in the evaluation of coronary artery disease, with the use of intravascular ultrasound (IVUS), with some good results if compared with the histological exams. Diethrich *et al.* [13] showed that, besides the coronary images, IVUS can also be applied to carotid ones to create a *VH*. Good correlation in the histological comparison was obtained. Nonetheless, IVUS is an invasive method.

The current most widely used methods for *VH* are based on defining ranges of pixel values for each tissue (*threshold methods*). The pixel is then classified based on its isolated value. Lal *et al.* [9], for instance, defined the following ranges of pixel values for each tissue: 2 (0–4) for blood, 12 (8–26) for lipid, 53 (41–76) for muscle, 172 (112–196) for fibrous tissue, and 221 (211–255) for calcium. Related methods also rely on the same idea [14], [15]. However, given the nature of ultrasound images (e.g., noise, differences in illumination) it is not straightforward to define such thresholds. Additionally, defining disjoint ranges of pixel values for each tissue can be impossible, as we show in the dataset used in the experiments herein. Fig. 4 shows the class overlap in the pixel intensity distribution of the dataset considered here.

In this context, this paper proposes a new approach for classifying tissues in ultrasound images at the pixel level, similarly to how it is done in VH IVUS, with the advantage that it is a totally noninvasive method. The proposed method depends only on the recorded B-mode image, so it could be performed after the patient examination. The method is based on the use of visual descriptors, including a new proposed descriptor that takes into account information from the neighborhood of a pixel, instead of considering isolated pixel values. The neighborhood of a pixel is also analyzed in multiple scales, capturing the pixel neighboring patterns, which are not captured by traditional threshold methods. The extracted features are used in a machine learning scheme for predicting the class of each pixel. On top of that, because of the difficulty of classifying each pixel due to the noisy nature of the problem, we also deploy a normalization method for stretching the dynamic range of the images.

We compare the proposed technique to existing ones in the literature showing its effectiveness. We also analyze the novel image descriptor used in the proposed method with respect to other descriptors of the literature. Additionally, we perform a histological specimen study that showed a good correlation with the classification results obtained by the proposed *VH* method.

II. Related Work

IVUS has been successfully validated in coronary arteries [16], [17]. However, IVUS uses a catheter with a small ultrasound probe and is an invasive method. In 2007, Diethrich *et al.* [13] evaluated *VH* IVUS imaging (VH IVUS) for carotid arteries and showed that VH IVUS is well correlated with the real histology.

As this paper focuses on noninvasive solutions, in this section, we present only non-invasive methods for tissue and carotid classification.

To solve the main problem of plaque characterization, the methods in the literature can be grouped into two categories: 1) methods that study the plaque using global information, i.e., based on the whole plaque image, for further classifying it as symptomatic or asymptomatic, and 2) methods that classify individual pixels, which can then be used to help in the diagnosis and to create the *VH*. The method proposed in this paper falls into the latter group.

A. Global Classification Methods

Mougiakakou *et al.* [18] tackled the problem of classifying symptomatic and asymptomatic plaques. The authors used a dataset of 54 symptomatic and asymptomatic images to train a binary classifier for the target problem. The feature extraction was done with the Laws' texture energy, resulting in 99 features that were reduced to 21 with analysis of variance. A neural network was trained with a genetic algorithm that selected the most robust features to train. This method has shown a classification accuracy of 95%.

Acharya *et al.* [19] used a dataset of 160 plaques (110 asymptomatic and 50 symptomatic) and, in all of them, 36 types of features were extracted using local binary pattern (LBP), Fuzzy Gray Level Co-occurrence Matrix, higher order spectrum features, and others. The classification was made using probabilistic neural networks, decision trees, and support vector machines (SVM). The best result reported was 90.6% of classification accuracy.

B. Pixel-Level Classification Methods

There are several methods in the literature aiming at classifying each pixel in an atherosclerotic plaque as one type of tissue, while requiring only data from B-mode ultrasound to do so. Almost all of them use the GSM value of the pixel for the classification.

Sztajzel *et al.* [20] proposed the stratified GSM and evaluated its use with color mapping to predict plaque histology. The stratified GSM is used to obtain a GSM value for each group of pixels. Each group of pixels is defined according to its pixel's

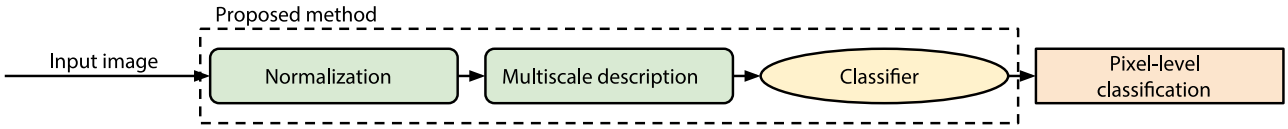

Fig. 1. Overview of the proposed method. The proposed method performs a normalization step followed by a multiscale description and uses a classifier for pixel classification. The multiscale description step is applied to each pixel of the image by considering its neighbors in several scales. Based on the pixel description, a classifier is trained to differentiate among the possible types of tissue.

distance to the plaque surface: at each millimeter increment, a new group is defined. The authors also used color mapping to predict plaque histology. Pixels were classified into three different colors: those with a gray-scale value under 50 as red, between 50 and 80 as yellow, and above 80 as green. The predominant color is then obtained with the aid of some strata of the stratified GSM compared to the total GSM value (GSM of the whole plaque). The authors demonstrated with real histopathological exams that analysis of color predominance with stratified GSM are correlated with the instability of the plaque. Among the conclusions, one of them is that unstable plaques have more of the red (0–50) color. In their work, the pixels are not individually classified as the possible tissue. Instead, the whole plaque is classified as either symptomatic or asymptomatic.

Lal *et al.* [9] analyzed several tissue images to calculate the GSM of each specific tissue. The values obtained with the mean of medians and the range of the median of all the images were 2 (0–4) for blood, 12 (8–26) for lipid, 53 (41–76) for muscle, 172 (112–196) for fibrous tissue, and 221 (211–255) for calcium. With those values, the atheroma plaque images were colored so that each color represented a tissue. It was observed that symptomatic plaques contain a high quantity of blood and fat while asymptomatic plaques contain more calcium.

Similar to Lal *et al.* [9], Madycki *et al.* [15] proposed a range of values to classify the pixel into one kind of tissue. These values (0–9 for blood, 10–31 for lipid, 32–74 for muscle, 75–111 for fibrous tissue, and 112–255 for calcified tissue), in contrast with Lal *et al.* [9], do not have any gaps, therefore no pixel would be predicted as unknown. The study showed that this analysis is more precise to predict the degree of microembolism, as compared to the normal GSM of the plaque.

One of the most recent articles for classifying each pixel of the plaque as a tissue is the one proposed by Hashimoto *et al.* [14]. The authors compared carotid endarterectomy specimens and the GSM values of known tissues on ultrasound B-mode images. The range of GSM values for each tissue was 0–24 for blood, 25–68 for lipid, 69–225 for muscle/fibrous tissues, and 226–255 for calcium. According to the authors, using the distribution of the pixel intensities in the plaque, was possible to predict the best diagnostic for the patient, i.e., whether the plaque was asymptomatic or if intervention was needed.

All those works showed the importance of discovering the plaque composition and how it can help physicians in routine exams. It is important to mention that all aforementioned *threshold-based methods* are computationally efficient, since they are based only on the isolated pixel value. For the same reason, however, the classification could be highly affected by the inherent noise of the image acquisition process, small changes of illumination, as well as inexperience of the DU operator. On top of that defining disjoint ranges for the pixel values can be impossible, as we show in the experiments (see Fig. 4). With this in mind, we propose a method that can be more robust to those problems and also sufficiently general to solve the tissue classification problem.

III. Proposed Method

There is a limited number of works in the literature focused on the tissue classification problem at the pixel level in DU images. Most of the methods explained in the previous section, the *threshold-based methods*, have a similar approach: classification of pixels based on their intensity. We have observed, however, that this classification methodology is not completely well suited when extended for classifying atheroma plaques, since it presents several problems regarding: noise, illumination changes, motion changes during capture, etc. In this sense, we devise an alternative way for approaching the problem. Our method is based on the premise that one pixel has a high probability of following its neighborhood pixel properties (spatial coherence), therefore, it is likely that if a small neighborhood around a pixel contains only blood pixels, then this pixel will probably be of the class blood. In addition, our characterization methodology also considers consistency across different scales. The rationale is that a pixel of a given tissue in one scale will be consistent across scales when incorporating more neighbors.

In Fig. 1, we present an overview of the proposed method. The main steps of the proposed method are the normalization of the input image, the description of each pixel based on its neighbors by considering multiple scales, and the use of a classifier for pixel classification.

We initially normalize the images to become more invariant to different ultrasound configurations, and to stretch the dynamic range of the images (see Section III-A). In the multiscale description step, we characterize each region of interest (RoI) (the region around the pixel under consideration) using a visual descriptor (e.g., statistical, gradient, and texture properties). This description considers multiple scales, i.e., considering different sizes of the RoI around the pixel (see Section III-B). Once a description is obtained for each training pixel, we train a classifier for pixel-level classification. In this study, we use the

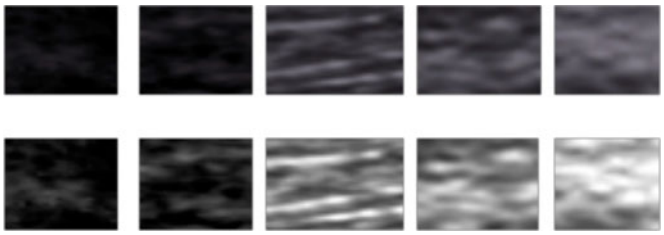

Fig. 2. Normalization of ultrasound tissue images. Nonnormalized (top) and normalized (bottom) blood, lipids, muscle, fibrous, and calcium examples, respectively.

multiclass-from-binary SVM using the pairwise approach [21] (see Section III-C).

In the remainder of this section, we present further details on the steps of the proposed method.

A. Normalization

Ultrasound image normalization is the main step to make possible the comparison of information obtained from images captured from different instrumental settings, scanners, operators, and any other factors that may change the image. Indeed, many studies [22], [23] demonstrate the importance of normalization methods for carotid plaque images. In this paper, we used the method proposed in [24].

The normalization of all the tissues of a person is made using the person's carotid artery image as a reference. The blood and the adventitia wall were chosen to take the values 0 and 190, respectively, and the other pixels in the image were linearly scaled according to these values. Thus, for each person, all the tissue images were normalized from the blood vessel image. Fig. 2 depicts some examples of tissue images before and after the normalization.

If no access to the carotid image of a person is possible or even the manual demarcation of the blood and adventitia wall, an automatic normalization process might be accomplished based on the *a priori* knowledge of blood and adventitia wall of some training images.

B. Multiscale Description

We are interested in classifying each pixel separately, hence all descriptors must take into account only a certain RoI around the pixel. We define a RoI of a pixel as a window of size $n \times n$ pixels (where n is an odd number) centered around the pixel currently in analysis. The image is then scanned using a sliding window approach. For describing each window position, i.e., describing each RoI, we evaluated different features, each one aiming at capturing different image properties.

The scale of the scanned images eventually varies depending on how the depth adjustment of the exam is performed. In order to guarantee a greater steadiness, multiscale windows/RoIs are used with all the types of descriptors tested (statistical-, gradient-, and texture-based). We preferred to use small RoIs (5×5 up to 13×13) in order to evaluate a pixel only in terms of its local information, assuring we do not break the pixel's spatial coherence assumption. For each descriptor, we calculated the feature vector on different window sizes and then combined them. The combination of feature vectors in each window size was made by concatenating the resulting feature vector of each scale.

TABLE I
DETAILS OF THE STATISTICS USED FOR SMD

Statistics	Symbol	Formula
Mean	$\bar{p}$	$\frac{1}{N}\sum_{j=1}^{N} p_j$
Standard Deviation	σ	$\sqrt{\frac{1}{N-1}\sum_{j=1}^{N}(p_j-\bar{p})^2}$
Skewness	γ_1	$\frac{1}{N}\sum_{j=1}^{N}\left[\frac{p_j-\bar{p}}{\sigma}\right]^3$
Kurtosis	γ_2	$\left(\frac{1}{N}\sum_{j=1}^{N}\left[\frac{p_j-\bar{p}}{\sigma}\right]\right)^4 - 3$
Median	m	The middle value of an ordered distribution of $p_1, p_2, \ldots, p_N$.
Entropy	$H(I)$	$-\sum_{j=1}^{N} p(p_j)\log_2 p(p_j)$, where $p(p_j)$ is the probability of occurrence of p_j
Range	R	$\max(p_j) - \min(p_j)$

Later, we present details of the descriptors used in this study: statistical moments descriptor (SMD), histograms of oriented gradients (HOG), LBP, statistical analysis of structural information (SASI), quantized compound change histogram (QCCH), Unser, and Steerable Invariant Descriptor (SID).

1) SMD: The *SMD*, which is a proposal of this study, discriminates the different forms of pixel intensity distribution within a RoI. Formally, we compute the following statistics: mean, standard deviation, skewness, kurtosis, median, entropy, and range (see Table I). The final feature vector f comprises a total of seven features, i.e., $f \in \Re^7$. For the definitions in Table I, p_j is the value of a pixel in the RoI, where j goes from 1 to N and $N = n^2$.

2) HOG: Visually, we can notice that certain tissues, such as muscle and fibrous, have some sort of orientation in lines. Therefore, we employed the *HOG* [25] to analyze the gradient orientations of the RoI. HOG works as follows. Let W be the RoI under consideration. Filtering W with the kernels $[-1\ 0\ 1]$ and $[-1\ 0\ 1]^T$, two filtered images W_x and W_y are obtained, which can be interpreted as the horizontal and vertical variations of the pixel, respectively. For each pixel in the RoI, the gradient orientation (Θ_W) and magnitude are calculated, which means that we obtain, respectively, the *direction* and the *intensity/magnitude* in which the pixel varies the most in relation to its neighbors. Then, the orientation values of the RoI are uniformly quantized into B intervals (usually $B = 9$) and a histogram of orientations is computed, generating a feature vector with dimensionality B for each RoI. The magnitude is the vote weight of the histogram.

3) LBP: An algorithm commonly used in the literature for describing textures with low gray-scale variance is the LBP [26].

The algorithm works as follows. For each pixel p in the RoI, we compare its value with the values of the eight-neighborhood pixels in a specific sequence. If the value of p is greater than or equal to the value of its kth neighbor, we write 1 in the kth position of an 8-bit binary string, otherwise, we write 0. Those 8 bits represent a decimal number between 0 and 255 that are the LBP code of p. Then, we create a histogram of 256 bins with the LBP codes of all pixels in the RoI and use it as the RoI feature vector.

4) SASI: The SASI [27] descriptor encodes structural properties of textures. Its feature extraction algorithm scans the RoI using windows with multiple resolutions and orientations. The first step of the algorithm is the choice of the window sizes and orientations to be used. Each window can be covered in different ways that are determined by vectors called *lag vectors* of value $K = \lceil (S/4) + 1 \rceil$, where S is the width of the window in pixels. For each window, the algorithm runs through the RoI calculating a value of autocorrelation considering different directions. At the end, the RoI has a set of autocorrelation values for each direction and each window. The mean and standard deviation of each set of autocorrelation values are calculated and stored sequentially as a feature vector. Then, the feature vector is normalized based on the mean and standard deviation of all autocorrelation values of all sets.

5) QCCH: The quantized compound change histogram descriptor [28] encodes the texture information by considering the relation of a pixel and its neighbors. The extraction algorithm defines a squared neighborhood (window) $N_r(i, j)$ of radius r centered at pixel (i, j) in the RoI. The RoI is scanned and, for each window position, the average gray value is computed. From this resulting image, four variation rates (horizontal, vertical, diagonal, and antidiagonal) are computed for each pixel and the average of those four values is calculated. The obtained average values are nonuniformly quantized into 40 bins and their histogram is used as feature vector.

6) Unser: The Unser [29] descriptor was proposed with the goal of reducing the complexity of the gray-level co-occurrence matrix (GCOM or GLCM). The main idea is to calculate two histograms, one of the sum and other of the differences between pixels. The algorithm initially defines an angle a to be considered between the neighboring pixels. The RoI is scanned and, for each angle a defined, each histogram is incremented independently: $H_{\text{sum}}[a][f(p) + f(q)]$ and $H_{\text{diff}}[a][f(p) - f(q)]$, where $f(x)$ is the pixel value, p is the current pixel, and q is its neighbor with an angle a in relation to p. Then, the histograms are normalized by the number of pixels in the RoI. From the histograms, it is possible to extract measures in a similar way to the co-occurrence matrix, which are: mean, standard deviation, correlation, energy, entropy, and homogeneity. Those measures compose the RoI feature vector.

7) SID: In the steerable pyramid decomposition descriptor [30], the input RoI is initially decomposed into two sub-bands using a high-pass filter and low-pass filter. Next, the resulting of the low-pass filter is recursively decomposed into K sub-bands by band-pass filters and one sub-band for a low-pass filter. Each step of the recursion captures different directional information in a given scale. The median and the standard deviation of each sub-band are stored as feature vector.

C. Classification

Our goal is to classify each pixel (represented by its feature vector obtained from its RoI) of an ultrasound image into one of the known tissue classes of interest (blood, fat, muscle, fibrous, and calcium). To do this, we need a supervised learning model that can learn from the data and then predict the correct class.

In this study, we used SVM [21] for the classification step. The SVM classifier constructs a hyperplane in the feature space, maximizing the margin distance between the positive and negative class samples. For the multiclass problem, one SVM classifier can be computed for all pairwise combinations of n classes. Therefore, for $n = 5$, we use $\frac{n(n-1)}{2} = 10$ binary SVM classifiers to decide among classes. The results of all those classifiers are combined by a majority voting scheme, yielding a final decision [21]. Although we have used SVM in this study, any other multiclass classifier would be applicable.

IV. Experiments

In this section, we present the experiments performed to evaluate the proposed method. Fig. 3 depicts the whole experimental procedure. We use two different datasets. Dataset 1 consists of tissue images, in which every pixel of each image has the same label, i.e., the entire images are of the same tissue. Dataset 2 comprises carotid images whereby different tissues can appear in the same image. For the second dataset, in addition to the pixel classification (VH), we have also conducted the actual histology of each image, thus, we could evaluate the correlation between the VH using our proposed method and the real histology.

We start by describing the datasets and the protocol used for the experiments. Then, we present the experimental results for each dataset.

A. Datasets

Ethics Statement: This study was approved by the Ethics Committee of the University of Campinas receiving the number of identification 277.976 in May 14, 2013. All participants provided written informed consent.

1) Dataset 1: Tissue Ultrasound Images: In order to extract the information of each tissue, we built a dataset with images of specific body parts that could represent a given tissue. The body parts chosen to be scanned were the cervical carotid artery (intravascular blood), the abdominal wall near the umbilicus (subcutaneous fat), the anterior aspect of the arm (biceps brachii for muscle), the posterior aspect of the lower leg (calcanean tendon for fibrous tissue), and the anterior aspect of the leg (tibia bone shaft for calcium). As we capture images of specific regions of the body that have high prevalence of one particular tissue, we made the assumption (verified by a specialist) that all the pixels in such images represent that tissue. The images were cropped to leave only the chosen tissue in the final image. The crop was made with the aid of a specialist, obtaining images with at most 150×150 pixels, in which most of the image contains the desired tissue.

All ultrasound images were obtained using an Acuson X300 equipment (SIEMENS AG, Munich, Germany), with a VF10-5 linear array transducer (set at 6.2 MHz) operating in B mode. We

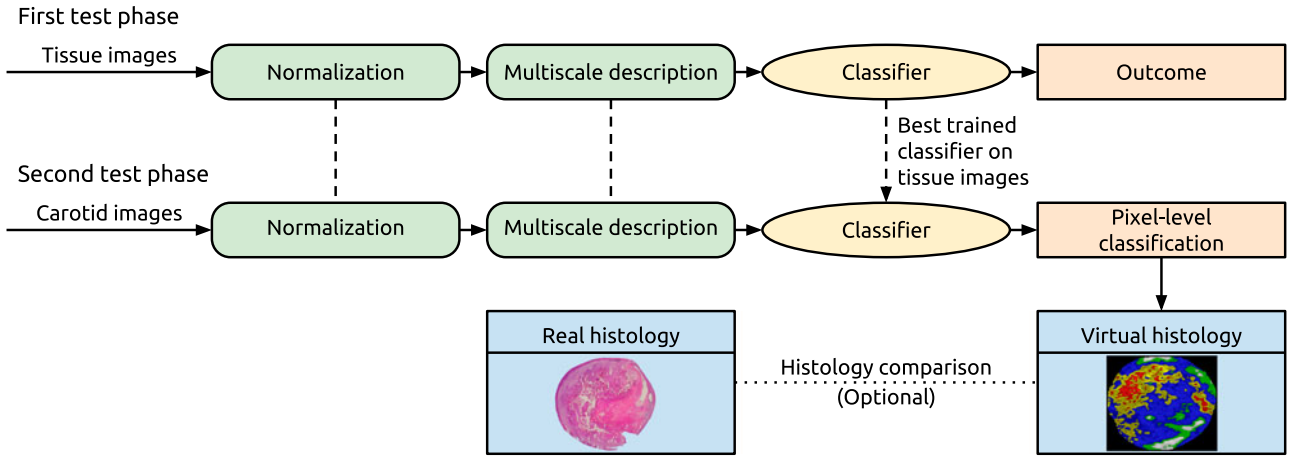

Fig. 3. Evaluation of the method. The tests using the proposed method are performed in two phases. In the first test phase, we use tissue images in which we assume that all pixels in each image are of the same tissue. In the second test phase, the best classifier obtained in the first phase is used in the carotid ultrasound images, generating the plaque coloring and consequently, the *VH*. Finally, we may compare virtual and real histology. The dash lines indicate that the linked boxes have the same configuration.

scanned 12 healthy volunteers (seven women and five men) and obtained five images of each tissue (five tissues in total) in two transducer orientations (transversal and longitudinal), totaling 50 images/person and 600 images in the whole dataset.

Due to the differences between the longitudinal and the transversal images, we have worked with them separately. For this paper, we mainly work with the images in the transversal way because the histological images are sliced in this direction.

2) Dataset 2: Carotid Ultrasound and Histological Images: To measure the effect of our algorithm in a practical setup, a second dataset of carotid images, scanned in the transversal plane, was obtained immediately before patients were submitted to carotid endarterectomy surgery. The images were obtained by a GE LOGIQ S6 device (Tokyo, Japan) with a 10L linear array transducer (10 MHz) in B mode. Six plaques of patients were analyzed. The patients that underwent surgery are symptomatic and have from 50 to 80 years old. The atheroma specimens were fixed in formaldehyde, sliced, and processed in the Pathology Department, University of Campinas. To guarantee that the ultrasound image and the histological specimen were compared in same localization, the ultrasound images were scanned in the location with maximum stenosis, and the plaques were sliced in the same place.

B. Experimental Protocol

We adopted the following protocol for measuring the accuracy of the methods: first, we separated the tissue images (Dataset 1) into two groups with the same size, one to train the classifier and find its parameters and the other to evaluate its performance. The same number of pixels were analyzed in the five classes to avoid unbalanced training. Also, there were no images of the same patient in both groups. This way, we guarantee that images of a person in the testing set were never seen by the classifier during training, which better simulates a real scenario of use.

Our protocol classifies pixels independently of movement. That is, the classification is performed frame-wise, not considering transitions between frames when scanning tissues. The analysis is made on a single frame that is acquired in a fraction of the cardiac cycle. The training was made in tissues that present no movement (muscle, calcium, fat, etc.).

For each scale (RoI size from 5 to 13), we randomly sampled 2000 RoIs (400 from each tissue class) for the training dataset. From this training set, we extracted the feature vectors using each descriptor described in Section III-B and trained an SVM multiclass classifier, performing grid search in the parameters C and γ [21]. Then, 5000 RoIs (1000 from each class) from the test partition were selected to measure the accuracy of the classifiers for each scale analyzed. The whole procedure was repeated ten times with different random partitions and samples for training and testing. Based on the ten runs, confidence intervals (95% of confidence) were computed. For statistical analysis, we used the Wilcoxon signed-rank test.

As ultrasound images have an inherent multiplicative noise called speckle, we also repeated the experiments for the images after noise reduction. On the one hand, the speckle noise itself makes the visual analysis difficult. On the other hand, the speckle itself could provide important information about the image and its tissues, depending on the extraction method used [31]. There are many filters that can be used to perform speckle noise reduction in ultrasound images [32]. We use the median filter, which replaces the center value of a 7×7-window by the median of the pixel values in such window. In the experiments, we compare the results in the original and in the filtered images.

To measure the accuracy in Dataset 1, we used the ratio of correctly classified pixels N_c in relation to the number of pixels in the image N. Formally, $Acc = N_c/N$.

Once we evaluated all descriptors with multiscale configurations in Dataset 1, we obtain the best trained classifier. Using this classifier, we performed the tests on the dataset of carotid images (Dataset 2) aiming at comparing the *VH* with the real one. This scheme is depicted in Fig. 3, where the first test phase refers to the tests performed on Dataset 1 of tissue images and in the second test phase the tests are performed on Dataset 2 of carotid images.

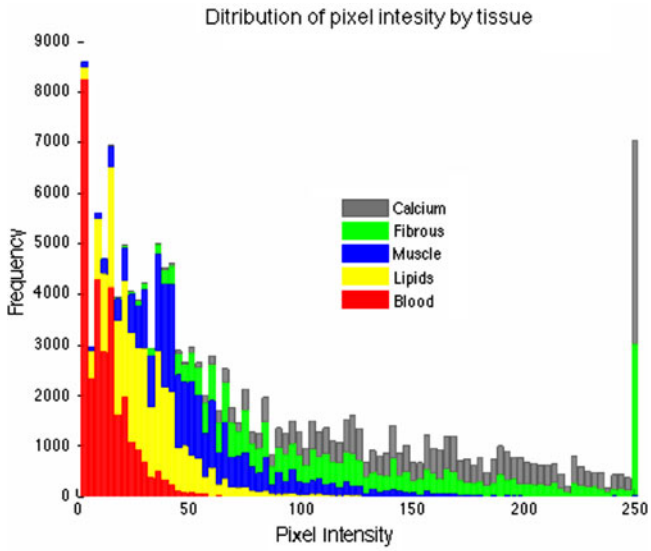

Fig. 4. Pixel Intensity distribution of the tissues. This histogram shows pixel-value frequencies for each type of tissue (Blood, Lipids, Muscle, Fibrous, and Calcium). Note the difficulty for establishing one threshold for classifying each type of tissue (there are several class overlaps).

C. First Test Phase: Experiments With Tissue Image Set

In this section, we describe experiments using the tissue images (Dataset 1). We first compare the image descriptors in a single-scale approach (see Section IV-C2). We also evaluate the descriptors by combining their results using a fusion scheme based on SVM scores. We perform a deeper evaluation of the best descriptor found (the SMD), considering the multiscale analysis (see Section IV-C3). Once we obtained the best combination of descriptor and multiscale configuration, we also evaluate the impact of the training set size (see Section IV-C4). Finally, we show a confusion matrix based on the possible classes of tissues, i.e., blood, lipids, muscle, fibrous, and calcium (see Section IV-C5).

1) Baseline: Threshold-Based Methods: To compare the different characterization approaches, we first define our baseline. In most related work, as presented in the Section II, the pixel classification is based only on its intensity value, which is defined according to prespecified ranges.

The problem in those approaches is that different sets of images may produce very different range values. Instead of using one of the threshold values proposed in the literature, which could have some bias due to the dataset, we decided to calculate the best threshold that could separate the tissues in our dataset. Using 30 000 pixels from the normalized tissue dataset (Dataset 1), we calculated a histogram of pixel intensities (from 0 to 255) for each type of tissue, with the same number of pixels per tissue. This procedure can be seen as a generalization of the aforementioned *threshold-based methods* [9], [14], [15] as it calculates the best ranges based on the training dataset.

The intersection of the histograms can be seen in Fig. 4. It is easy to see that the pixel values of different tissues are superposed. For instance, fibrous and calcium have similar frequencies in almost all bins. To find the best range of values, we used a brute-force approach calculating the accuracy for all possible different ranges. The best accuracy achieved was 54.01% with the following intervals for each tissue: blood (0-16), lipid (17-37), muscle (38-83), fibrous tissue (84-160), and calcified tissue (161-255).

2) Descriptor Comparison: In Table II, we show the accuracies of each descriptor in isolated scales, i.e., when using only one RoI size during feature extraction. We are showing only the scales in which descriptors presented better accuracies. We can see that proposed SMD achieves the highest accuracy rates, which are in fact, far above the other descriptors.

For SMD, the smaller RoI (9×9) was only slightly worse than the larger one (13×13). A different behavior can be observed for the other descriptors, which were significantly worse with the smaller RoI. As most of the other descriptors are based on texture, having a larger neighborhood could provide more information about the pixel's local texture. On top of that, the other descriptors are usually not employed for pixel classification, being popularly used for image classification. SMD is invariant to rotation and small translation. Additionally, as many of the statistics used in SMD are not largely affected by the inclusion of possible outliers, its accuracy remains similar in both 13×13 and 9×9 RoI sizes.

Comparing the results for the original and the filtered images, we see that there is no statistical difference (Wilcoxon signed-rank test) for most of the descriptors. This result suggests that the descriptors were not affected by the speckle noise in the dataset.

Comparing the results with the baseline (54.01%), we can see that for the larger RoI (13×13), the descriptors SMD, SASI, QCCH, Unser, and SID are better than the baseline. For the smaller RoI, SMD is the only one to be better than the baseline. Those results illustrate that by considering the pixel neighborhood, we can obtain a better estimation of the pixel class, in comparison with the strategy that uses only the isolated pixel value.

It is worth noting that the baseline approach could "learn" from a larger portion of the training set in comparison to the other descriptors. For finding the thresholds for the baseline, 30 000 pixels were used, but for the other descriptors, only 2000 pixels were used for training the classifier. The possible effect of this difference would be an improvement in the accuracy of the baseline, but we could see that even with this advantage, the baseline was worse.

2) Descriptor Fusion: With the aim of improving the classifier's performance, we decided to evaluate the fusion of the predictions of all the descriptors. The fusion strategy is based on the classification probability function of the classifier (late fusion) [21], [33], which estimates the probability of a given test sample of belonging to each class. For each descriptor, we acquire the probability for each class. For the fusion of n descriptors, the final decision value for the class C is the sum of the n probability values for C. The class with the highest final decision value is chosen.

Even by considering all the possible combinations of descriptors/classifiers, we were not able to find a combination that was statistically better than the single proposed SMD (no

TABLE II
DESCRIPTOR ACCURACIES FOR ISOLATED SCALES (ROI SIZES 9 AND 13)

(a) RoI size 9×9			(b) RoI size 13×13		
Descriptor	**Original Images**	**Filtered Images**	**Descriptor**	**Original Images**	**Filtered Images**
SMD	**69.58 ± 0.94**	**69.23 ± 0.58**	**SMD**	**71.61 ± 0.66**	**71.24 ± 0.71**
HOG	47.66 ± 0.87	49.10 ± 0.64	HOG	50.47 ± 0.72	51.00 ± 0.79
LBP	34.23 ± 0.35	34.16 ± 0.37	LBP	48.59 ± 0.95	48.57 ± 1.18
SASI	35.16 ± 0.25	34.94 ± 0.58	SASI	55.39 ± 1.19	56.99 ± 0.57
QCCH	35.33 ± 0.43	34.72 ± 0.48	QCCH	63.75 ± 0.42	63.87 ± 0.50
Unser	37.35 ± 0.35	35.61 ± 0.41	Unser	66.34 ± 0.61	67.92 ± 0.65
SID	35.27 ± 0.62	35.01 ± 0.43	SID	64.95 ± 0.55	64.75 ± 0.58
Fusion	69.76 ± 0.74	69.55 ± 0.53	Fusion	71.26 ± 0.52	71.19 ± 0.31

We can see that SMD achieves the highest accuracies. Filtering images for noise reduction did not affect the results. Fusion was not effective as well.

TABLE III
ACCURACIES WHEN COMBINING MULTIPLE ROI SIZES FOR THE SMD

RoI sizes	**Accuracy**
7 and 9	70.39 ± 0.48
9 and 11	70.97 ± 0.37
11 and 13	71.88 ± 0.08
5, 7, and 9	70.14 ± 0.54
7, 9, and 11	70.09 ± 0.59
9, 11, and 13	**73.17 ± 0.38**
5, 7, 9, and 11	69.87 ± 0.54
7, 9, 11, and 13	72.02 ± 0.40

We can note an improvement in accuracy with the multiscale description.

fusion). The best fusion results are in the last row of Table II. Therefore, we decided to focus the following experiments with the best image characterization form we explored: the SMD.

3) Multiscale Analysis: Here, we performed a deeper analysis using SMD, as it was the best descriptor found in the previous experiments. We first evaluate how the combination of multiple RoI sizes can improve the results. We then evaluate the impact of the training set size.

When the size of the analyzed RoI is larger, more neighborhood information is available. This explains why the 13×13 RoI size achieves better results than other smaller sizes. However, note that it is important to limit the RoI size to a maximum in order not to break the pixel's spatial coherence assumption. Some scales may contain important information that others could not capture because operators of the ultrasound machines might use different torque while capturing the images, therefore, we decided to combine feature vectors of different scales and evaluate the performance of SMD.

In Table III, we show some of the combinations and their accuracies. The best result was obtained with three scales of sizes 9, 11, and 13. This result was statistically different from the one using only one scale of size 13 (Wilcoxon signed-rank test, two tail, p-value $= 0.0013$), as well as from the best results with two and four scales (p-value $= 0.014$ and p-value $= 0.0098$, respectively). Therefore, the feature vectors with scales of sizes 9, 11, and 13 were chosen to be combined (concatenated) for generating an optimal descriptor: Multiscale Statistical Moments Descriptor (MS-SMD).

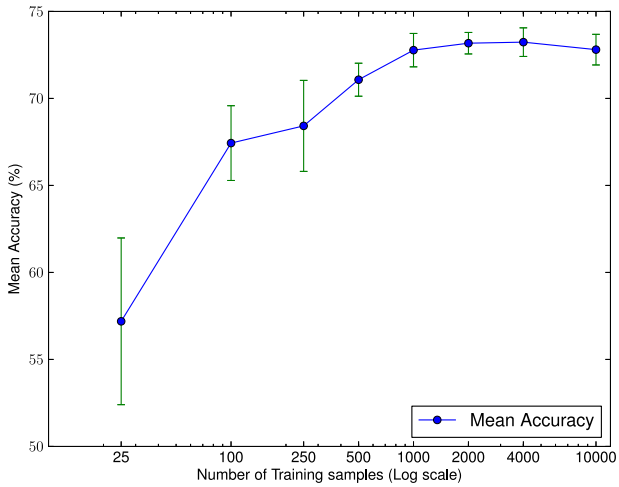

Fig. 5. Training-set size versus classification accuracy for MS-SMD. Performance stabilizes for 1,000+ training samples.

TABLE IV
CONFUSION MATRIX OF THE TISSUE CLASSIFICATION USING MS-SMD

	Blood	Lipids	Muscle	Fibrous	Calcium
Blood	**91.35**	8.62	0.02	0	0
Lipids	9.91	**79.63**	10.42	0.04	0
Muscle	0.23	16.03	**63.49**	18.83	1.42
Fibrous	0	0.01	12.26	**55.67**	32.05
Calcium	0	0	1.65	22.62	**75.71**

Rows represent the real class and columns, the predicted class. The main confusion is between fibrous and calcium tissues. However, for blood and lipids (the most important tissues when analyzing a carotid plaque image), the accuracies are the highest ones.

4) Training Set Size: Having set MS-SMD with the best scale as the combination of RoI sizes 9, 11, and 13, we then sought to determine if the size of the training set is relevant to the classification. To do so, we varied the number of RoIs to train from 25 to 10 000. For each training set size, we used the same protocol: randomly selected the same number of RoIs per tissue class for training, and testing with 5000 RoIs randomly selected from the images not used for training. The experiment was repeated ten times for each set size as well as in the previous sections.

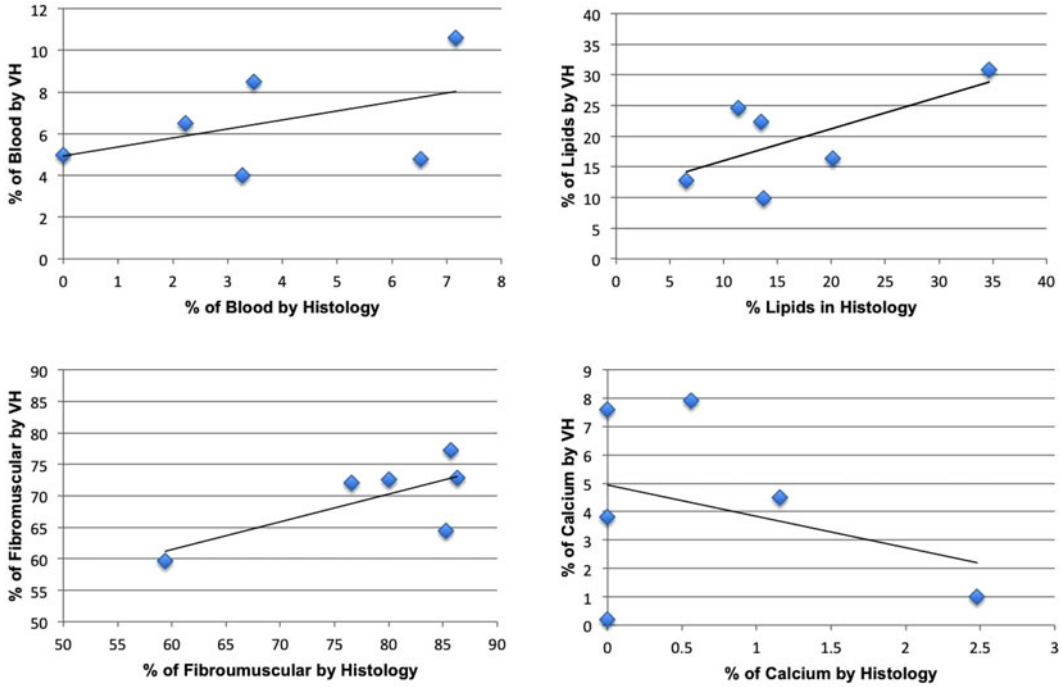

Fig. 6. Pearson correlation of each type of tissue. The correlation values for each tissue are: blood ($\rho = 0.456$), lipids ($\rho = 0.652$), fibromuscular ($\rho = 0.704$), and calcium ($\rho = -0.338$).

Fig. 5 depicts the results for different training sizes. The Wilcoxon signed-rank test was used for comparing the results using pairs of training set sizes, e.g., 25×100, 100×250, and so on. The classification accuracy stabilizes for 1000+ training samples, i.e., from this point on, there is no statistical significance in changing the training set sizes.

5) Results: The final classification approach obtained is the one using MS-SMD (multiple RoIs of sizes 9, 11, and 13) with 1000 training examples per class. Table IV shows the classification confusion matrix for MS-SMD. Note, for instance, that fibrous and calcium are tissues with a similar echogenicity, so even with the MS-SMD approach, we were not able to totally eliminate the misclassification between these two classes. Blood and lipids, that also have a similar echogenicity, were well separated, with a small classification error. Note also that for analyzing a carotid plaque image, the two most important tissues are blood and lipids and both have a good classification accuracy.

It is important to highlight that this is a pixel-level classification, therefore it is not easy to achieve a perfect classification score, since we are predicting the entire image pixel-wise. Another problem is that it is impossible to ensure that 100% of the training images really contain only pixels of the labeled tissues. For instance, in a tibia bone, almost all the tissue is calcium, however it may contain parts of other kinds of tissues. On top of that, there is the universal presence of *connective tissue* in nearly all tissues of the human body [34].

D. Second Test Phase: Experiments With Carotid and Histology Set

In this section, we describe the experiments on Dataset 2 related to the "second test phase" of Fig. 3, whereby we perform *VH* (see Section IV-D1) and then compare the results with the histology on real carotid plaques (see Section IV-D2).

1) Plaque Coloring: Beyond classifying each pixel individually, it is important to show the classification in a visible way in the plaque in order to facilitate the work of specialists. This is done with the concept of *VH*, in which each pixel of the ultrasound image is colored according to the tissue classification. In our case, blood is seen as red, lipids as yellow, muscles as blue, fibrous as green, and calcium as white. We also set the intensity of the color according to the degree of classifier certainty in the prediction (see Section III-C). For example, if the classifier points out with a higher probability that a pixel represents lipid, it will be colored with a brighter tone of yellow. Some of the plaques colored by our method are shown in Fig. 7.

2) Histology Comparative: After coloring the carotid plaque in the ultrasound images, we have performed the real histology in the same plaques. According to previous studies, the comparison of histological specimens to *in vivo* images (and even *in vitro* images) is a difficult methodological procedure [35]. When the atherosclerotic plaque is gathered at surgery and during the preparation of the slices for microscopy, a series of potential errors are introduced, as the loss of the true lumen, disruption of the integrity of the plaque, loss of constituents through the use of solvents, etc. Also, as specimens are imaged *ex-vivo* they do not represent the true aspect found *in vivo*, as the arterial wall is disrupted and the luminal arterial pressure is not present, rendering the conclusions not applicable to the real patient. Another important problem is the alteration in the gray-scale brightness after embedding the specimen in gel for ultrasound visualization or after preserving the specimen in formaldehyde, which potentially renders the pixel level evaluation not reliable. Therefore, we should take such experiments of histology comparison

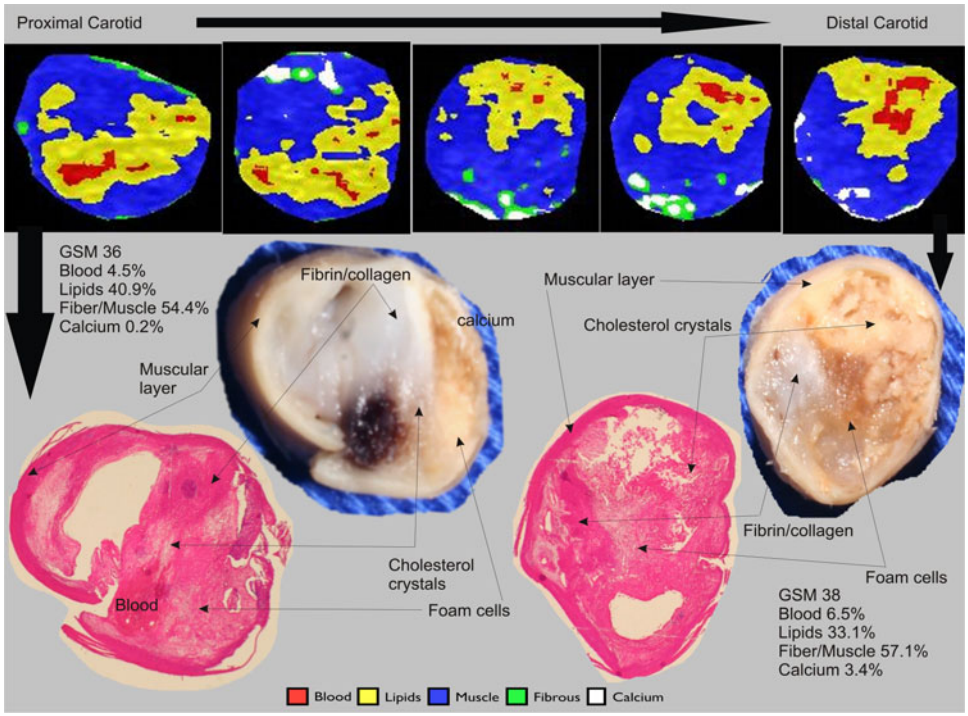

Fig. 7. Illustration of the comparative findings between *VH* of ultrasound images and macroscopic and microscopic pathological findings. The five images at the top correspond to five different ultrasound images of the atherosclerotic plaque *in vivo* from proximal to distal internal carotid starting at its bifurcation position from the carotid bulb. The distance from the first and last images is approximately 5 mm. It can be seen that the plaque content varies in quantity and position in the different slices. Below them, we have the macroscopic and microscopic images of the plaque cut at approximately the same first (top-left image) and last (top-right image) image slices. It can be observed that the lipid rich foam cells and the cholesterol crystal deposits produce a gross appearance with a brown–yellow color and appears in the *VH* images as yellow, and sometimes red, pixels. The fibrin/collagen deposits inside the plaque and the peripheral muscular layer appear white in the gross macroscopic view and as blue in the processed ultrasound images. The areas that lack structures in the microscopic images are due to imperfections in the surface of the gross specimen, an artifact error that is not found in the *in vivo* ultrasound images.

(real and virtual) as a complement to the experiments in the ultrasound tissue images.

All the histological slices were digitized and analyzed by an experienced pathologist. For each plaque, the histological components calcium, fibromuscular tissue, lipid core, and hemorrhage were selected and their area calculated. Then, for each type of tissue, we performed the Pearson correlation between the percentages in the real histology and *VH*. Blood is correlated with hemorrhage, lipids with lipids, calcium with calcium, and the compound fibrous and muscular tissue with fibromuscular tissue.

Fig. 6 depicts the plots of each type of tissue. In total, we analyzed six plaques of patients sent to surgery. Analyzing the results, we see that the results of the *VH* using the proposed method have similar percentage with histology exams for lipids and fibromuscular tissue. Blood also has a considerable similarity with the real histology exams.

In Fig. 7, we show a comparison of the carotid plaques colored by the proposed method (VH) and the real histology. The five images at the top correspond to five different ultrasound images of the atherosclerotic plaque *in vivo* from proximal to distal internal carotid starting at its bifurcation position from the carotid bulb. Below them, we have the macroscopic and microscopic images of the plaque cut at approximately the same first (top-left image) and last (top-right image) image slices. We can see, for instance, that the cholesterol crystal deposits produce a gross appearance with a brown–yellow color and this same region appears mostly in yellow (lipids) in the VH (results from our algorithm). The observation of the images presented at the many studies comparing imaging modalities to histological slices demonstrate that the plaque constituents are easily differentiated in both microscopic image and imaging modality, and we have had the same perception in the analyzed plaques, thus, far. More images comparing real and virtual histologies are available in the supplementary material along with this paper.

V. Conclusion and Future Work

In this paper, we presented a new method for classifying tissues in ultrasound imagery at the pixel level. The proposed method is basically composed of 1) image normalization, 2) multiscale description, and 3) classification. The main advantage of the proposed method is its capacity to consider pixel-neighboring patterns, which are not encoded by existing methods based on isolated pixel values. By also performing multiscale description of image regions, the method can deal with the intrinsic scale variability of the problem, as operators of ultrasound devices might use different torque while capturing the images.

The proposed method was experimentally evaluated in tissue and carotid images. A comparison with the state-of-the-art techniques for pixel-level classification [9], [14], [15] shows the quality of the proposed method and its advantages. The best threshold-based method (a generalization of [9], [14], [15] with thresholds dynamically calculated based on a training set) achieved about 54% of pixel classification accuracy, while the proposed method achieved more than 73%. Additionally, we also show in the paper the correlation of the VH obtained by the proposed method with the real histology.

A prototype implementing the proposed method is currently in use in the University of Campinas Hospital, Brazil, and the analyses conducted thus far indicate a good correlation with the real histology and more confident results than threshold-based methods. Obtaining better pixel-level tissue classification is paramount for further providing a better plaque classification in terms of asymptomatic or symptomatic, and so, for determining the risk of diseases for the patient.

The main limitation of the proposed method is the time required to perform the computation compared with the other methods in the literature. For instance, to classify and normalize a plaque region of around 100×100 pixels, the proposed method takes about 5–15 s, while threshold-based methods take about 1 s. The execution time was measured in Matlab, in a machine with 8 GB of RAM and 2.4-GHz Intel Core i5 processor.

Although validated in the problem of carotid plaque classification, we believe that the method could also be used in the classification of tissues in other types of ultrasound images as well, with some specific changes such as the normalization and scaling steps. The proposed technique may be expanded to evaluate venous disease, since the age of the intraluminal thrombus in deep venous thrombosis leads to changes in the image brightness. The classification of edema, including intracavity and parenchimatous fluids, may also benefit from such evaluations. In addition, there are studies evaluating the aspect of nodules in order to differentiate benign from malignant disease, these studies could also benefit from our solutions presented herein. In this way, we believe that the computational improvement of the classification methods for pixel brightness level is a promising technology with a long growing path ahead. On top of that, we believe that the proposed method is general enough to work with other kinds of medical images, as computational tomography, for instance. Tests in such applications and imaging technologies can be interesting as possible future work.

To improve and validate the robustness of the algorithm, one could train a specific classifier including images from many different types of ultrasound devices to account for the natural variability of acquisition devices. Furthermore, it is possible to create a specific tissue classifier for each person, with the person's own ultrasound images. This would eliminate the variation between patients and create a more accurate classifier to exam a patient. To make this possible, it would be necessary to scan a few images of specific tissues of this person (e.g., the abdominal wall near the umbilicus for subcutaneous fat, the anterior aspect of the arm for muscle, etc.), and then train the specific classifier to examine the carotid's image.

ACKNOWLEDGMENT

The authors would like to thank CNPq, CAPES, FAPESP (Grant 2012/20465-5) for the financial support. Part of the results presented in this work was obtained through the project "Pattern recognition and classification by feature engineering, *-fusion, open-set recognition, and metarecognition," supported in part by Samsung Eletrônica da Amazônia Ltda., in the framework of law No. 8,248/91.

REFERENCES

[1] Executive Committee for the Asymptomatic Carotid Atherosclerosis Study, "Endarterectomy for asymptomatic carotid artery stenosis," *J. Amer. Med. Assoc.*, vol. 273, no. 18, pp. 1421–1428, 1995.

[2] V. A. Mantese, C. H. Timaran, D. Chiu, R. J. Begg, and T. G. Brott, "The carotid revascularization endarterectomy versus stenting trial (crest): Stenting versus carotid endarterectomy for carotid disease," *Stroke*, vol. 41, no. 10 Suppl, pp. S31–S34, 2010.

[3] MRC Asymptomatic Carotid Surgery Trial (ACST) Collaborative Group, "Prevention of disabling and fatal strokes by successful carotid endarterectomy in patients without recent neurological symptoms: Randomised controlled trial," *Lancet*, vol. 363, no. 9420, pp. 1491–1502, 2004.

[4] P. Gao, Z.-Q. Chen, Y.-H. Bao, L.-Q. Jiao, and F. Ling, "Correlation between carotid intraplaque hemorrhage and clinical symptoms: Systematic review of observational studies," *Stroke*, vol. 38, no. 8, pp. 2382–2390, 2007.

[5] R. D. Malgor, E. A. Wood, O. A. Iavarone, and N. Labropoulos, "Stratifying risk: Asymptomatic carotid disease," *J. Vascular Brasileiro*, vol. 11, no. 01, pp. 43–52, 2012.

[6] L. Hermus, G. M. van Dam, and C. J. Zeebregts, "Advanced carotid plaque imaging," *Eur. J. Vascular Endovascular Surg.*, vol. 39, no. 2, pp. 125–133, 2010.

[7] N. Sakalihasan and J. B. Michel, "Functional imaging of atherosclerosis to advance vascular biology," *Eur. J. Vascular Endovascular Surg.*, vol. 37, no. 6, pp. 728–734, 2009.

[8] L. A. V. Baroncini, A. Pazin Filho, L. O. Murta Junior, A. R. Martins, S. G. Ramos, J. Cherri, and C. E. Piccinato, "Ultrasonic tissue characterization of vulnerable carotid plaque: Correlation between videodensitometric method and histological examination," *Cardiovascular Ultrasound*, vol. 4, no. 8, p. 32, 2006.

[9] B. K. Lal, R. W. Hobson, P. J. Pappas, R. Kubicka, M. Hameed, E. Y. Chakhtoura, Z. Jamil, F. T. Padberg, P. B. Haser, W. N. Durán, and E. Y. Chakhtura, "Pixel distribution analysis of b-mode ultrasound scan images predicts histologic features of atherosclerotic carotid plaques," *J. Vascular Surg.*, vol. 35, no. 6, pp. 1210–1217, 2002.

[10] A. D. Montauban van Swijndregt, H. R. Elbers, F. L. Moll, J. de Letter, and R. G. Ackerstaff, "Ultrasonographic characterization of carotid plaques," *Ultrasound Med. Biol.*, vol. 24, no. 4, pp. 489–493, 1998.

[11] N. El-Barghouty, G. Geroulakos, A. Nicolaides, A. Androulakis, and V. Bahal, "Computer-assisted carotid plaque characterisation," *Eur. J. Vascular Endovascular Surg.*, vol. 9, no. 4, pp. 389–393, 1995.

[12] S. K. Kakkos, J. M. Stevens, A. N. Nicolaides, E. Kyriacou, C. S. Pattichis, G. Geroulakos, and D. Thomas, "Texture analysis of ultrasonic images of symptomatic carotid plaques can identify those plaques associated with ipsilateral embolic brain infarction," *Eur. J. Vascular Endovascular Surg.*, vol. 33, no. 4, pp. 422–429, 2007.

[13] E. B. Diethrich, M. P. Margolis, D. B. Reid, A. Burke, V. Ramaiah, J. A. Rodriguez-Lopez, G. Wheatley, D. Olsen, and R. Virmani, "Virtual histology intravascular ultrasound assessment of carotid artery disease: the carotid artery plaque virtual histology evaluation (capital) study," *J. Endovascular Therapy*, vol. 14, no. 5, pp. 676–686, 2007.

[14] H. Hashimoto, M. Tagaya, H. Niki, and H. Etani, "Computer-assisted analysis of heterogeneity on b-mode imaging predicts instability of asymptomatic carotid plaque," *Cerebrovascular Diseases*, vol. 28, no. 4, pp. 357–364, 2009.

[15] G. Madycki, W. Staszkiewicz, and A. Gabrusiewicz, "Carotid plaque texture analysis can predict the incidence of silent brain infarcts among patients undergoing carotid endarterectomy," *Eur. J. Vascular Endovascular Surg.*, vol. 31, no. 4, pp. 373–380, 2006.

[16] A. Nair, B. D. Kuban, E. M. Tuzcu, P. Schoenhagen, S. E. Nissen, and D. G. Vince, "Coronary plaque classification with intravascular ultrasound radiofrequency data analysis," *Circulation*, vol. 106, no. 17, pp. 2200–2206, Oct. 2002.

[17] K. Nasu, E. Tsuchikane, O. Katoh, D. G. Vince, R. Virmani, J.-F. Surmely, A. Murata, Y. Takeda, T. Ito, M. Ehara, T. Matsubara, M. Terashima, and T. Suzuki, "Accuracy of in vivo coronary plaque morphology assessment: A validation study of in vivo virtual histology compared with in vitro histopathology," *J. Amer. College Cardiol.*, vol. 47, no. 12, pp. 2405–2412, Jun. 2006.

[18] S. G. R. Mougiakakou, S. Golemati, I. Gousias, A. N. Nicolaides, and K. S. Nikita, "Computer-aided diagnosis of carotid atherosclerosis based on ultrasound image statistics, lawst'exture and neural networks," *Ultrasound Med. Biol.*, vol. 33, no. 1, pp. 26–36, 2007.

[19] U. R. Acharya, M. M. R. Krishnan, S. V. Sree, J. M. Sanches, S. Shafique, A. Nicolaides, L. M. Pedro, and J. S. Suri, "Plaque tissue characterization and classification in ultrasound carotid scans: A paradigm for vascular feature amalgamation," *IEEE Trans. Instrumentation Meas.*, vol. 62, no. 2, pp. 392–400, Feb. 2013.

[20] R. Sztajzel, S. Momjian, I. Momjian-Mayor, N. Murith, K. Djebaili, G. Boissard, M. Comelli, and G. Pizolatto, "Stratified gray-scale median analysis and color mapping of the carotid plaque correlation with endarterectomy specimen histology of 28 patients," *Stroke*, vol. 36, no. 4, pp. 741–745, 2005.

[21] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Trans. Intell. Syst. Technol.*, vol. 2, no. 3, pp. 27:1–27:27, 2011.

[22] M. M. Sabetai, T. J. Tegos, A. N. Nicolaides, S. Dhanjil, G. J. Pare, and J. M. Stevens, "Reproducibility of computer-quantified carotid plaque echogenicity: Can we overcome the subjectivity?" *Stroke*, vol. 31, no. 9, pp. 2189–2196, 2000.

[23] T. Tegos, M. Sabetai, A. Nicolaides, G. Pare, T. Elatrozy, S. Dhanjil, and M. Griffin, "Comparability of the ultrasonic tissue characteristics of carotid plaques," *J. Ultrasound Med.*, vol. 19, no. 6, pp. 399–407, 2000.

[24] T. Elatrozy, A. Nicolaides, T. Tegos, A. Zarka, M. Griffin, and M. Sabetai, "The effect of b-mode ultrasonic image standardisation on the echodensity of symptomatic and asymptomatic carotid bifurcation plaques," *Int. Angiol.*, vol. 17, no. 3, pp. 179–186, 1998.

[25] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Proc. Conf. Comput. Vision Pattern Recog.*, 2005, vol. 1, pp. 886–893.

[26] T. Ojala, M. Pietikainen, and D. Harwood, "Performance evaluation of texture measures with classification based on kullback discrimination of distributions," in *Proc. Int. Conf. Pattern Recog.*, 1994, vol. 1, pp. 582–585.

[27] A. Çarkacioğlu and F. Yarman-Vural, "SASI: A generic texture descriptor for image retrieval," *Pattern Recog.*, vol. 36, no. 11, pp. 2615–2633, 2003.

[28] C.-B. Huang, and Q. Liu, "An orientation independent texture descriptor for image retrieval," in *Proc. Int. Conf. Commun., Circuits Systems*, 2007, pp. 772–776.

[29] M. Unser, "Sum and difference histograms for texture classification," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-1, no. 1, pp. 118–125, Jan. 1986.

[30] J. A. Montoya-Zegarra, N. J. Leite, and R. D. S. Torres, "Rotation-invariant and scale-invariant steerable pyramid decomposition for texture image retrieval," in *Proc. Brazilian Symp. Comput. Graphics Image Process*, 2007, pp. 121–128.

[31] J. Noble, "Ultrasound image segmentation and tissue characterization," *J. Eng. Med.*, vol. 224, no. 2, pp. 307–316, 2010.

[32] C. P. Loizou, C. S. Pattichis, C. I. Christodoulou, R. S. Istepanian, M. Pantziaris, and A. Nicolaides, "Comparative evaluation of despeckle filtering in ultrasound imaging of the carotid artery," *IEEE Trans. Ultrasonics, Ferroelectrics Frequency Control*, vol. 52, no. 10, pp. 1653–1669, Oct. 2005.

[33] J. C. Platt, "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods," in *Advances In Large-Margin Classifiers*. Cambridge, MA, USA: MIT Press, 1999, pp. 61–74.

[34] A. Tsamis, J. T. Krawiec, and D. A. Vorp, "Elastin and collagen fibre microstructure of the human aorta in ageing and disease: A review," *J. Royal Soc. Interface*, vol. 10, no. 83, p. 20121004, 2013.

[35] J. K. Lovett, J. N. E. Redgrave, and P. M. Rothwell, "A critical appraisal of the performance, reporting, and interpretation of studies comparing carotid plaque imaging with histology," *Stroke*, vol. 36, no. 5, pp. 1085–1091, 2005.

Authors' photographs and biographies not available at the time of publication.