

Diagnosis of Alzheimer’s disease using 3D Local Binary Patterns

Pedro Morgado, Margarida Silveira, Jorge S. Marques

Instituto Superior Técnico

Instituto de Sistemas e Robótica

Lisbon, Portugal

pedro.morgado@ist.utl.pt, msilveira@isr.ist.utl.pt, jsm@isr.ist.utl.pt

Abstract

In the last decade, the computerized diagnosis of Alzheimer’s disease (AD) and Mild Cognitive Impairment (MCI) using the information provided by different neuroimaging techniques has been extensively studied. However, the texture of such neuroimages has been little explored. In this work, both diagnosis were conducted based solely on the texture of FDG-PET images, which was extracted using a novel three-dimensional extension of the well-known two-dimensional texture descriptor Local Binary Patterns (LBP). In LBPs, the concepts of uniformity and rotation invariance are of fundamental importance. We show that the proposed approach, unlike other 3D extensions found in the literature, closely replicates these concepts, as originally proposed in the 2D setting. Experimental results showed that the new 3D LBP version is able to enhance the generalization ability of the diagnostic system and also that the texture of FDG-PET scans contains distinctive information about the presence of both AD and MCI.

Keywords

Alzheimer’s Disease, Mild Cognitive Impairment, Computer-Aided Diagnosis, Texture Extraction, 3D Local Binary Patterns

I. INTRODUCTION

Alzheimer’s disease is a neurological disorder characterized by a severe loss of memory and cognitive abilities, such as planing and reasoning, and by the gradual onset and worsening

of symptoms [1]. It affects mostly people over 65 years old [2] and the average lifespan after diagnosis is about 8 years [3], although with a large variability. Moreover, AD patients are not the only ones affected. As the disease progresses, individuals become completely dependent on others (typically family members) even for the most basic daily tasks. All such unpaid hours of care and medical costs make AD a very expensive disease [4]. Even more alarming is the fact that its incidence rate and, consequently, its economic burden and the number of deaths AD related are still increasing due to the demographic ageing and population growth.

Although AD remains incurable to the present date, early detection is very important for an effective treatment able to slow down the progression of symptoms, improve life quality and extend life expectancy. A disorder that is typically associated with this early stage is Mild Cognitive Impairment [5]. The diagnosis of this early state is, however, difficult to perform due to two main reasons. First, because the onset of AD is often confused with the natural ageing process or linked to stress [6] and, second, because no characteristic pattern of brain degeneration is well-defined which adds uncertainty to the diagnosis.

Nevertheless, computer-aided diagnosis (CAD) of AD and MCI has revealed a great potential to improve diagnostic accuracy. In fact, not only a large number of CAD systems have already achieved performances at least as good as the ones attained by expert physicians [7], but also there is a promising ongoing research. This is a consequence of the high computing power available to these systems, which allows a much more sensitive analysis of the available information. In addition, CAD systems have the advantage of being less subjective and prone to error than human examination.

Positron Emission Tomography (PET) is a medical imaging technique that is often used by physicians to help in the diagnostic procedure, and which is able to track the consumption of a molecule injected into the body and known as tracer. Thus, when a tomography is performed on the brain and FDG (a glucose analogue) is used as the tracer, this technique is able to estimate the cerebral metabolic rate for glucose (CMR_{glc}) which is linked with brain activity and, therefore, with AD [8]. In fact, quantitative studies performed on FDG-PET images have found significant reductions of the CMR_{glc} in several regions of an affected brain [8]. Due to the aforementioned reasons, FDG-PET scans are often used as the main source of information of CAD systems.

Texture analysis has previously showed its utility as a tool to increase the information retrieved from medical images. It has been used, for instance, to distinguish healthy from damaged tissue in different organs and even to segment certain anatomical structures [9], but it has been little explored for the AD diagnosis. Herein, we will describe a system for the CAD of both AD and MCI exploiting the textural information of FDG-PET images. More precisely, we will use the well-known descriptor Local Binary Pattern for texture extraction and a Support Vector Machine for learning purposes. In addition, due to the three-dimensional nature of the biomarker in use, we will propose a novel extension of LBPs to 3D data.

The remaining of this paper is organized as follows. First, previous relevant works regarding both AD diagnosis and LBPs are reviewed in section II. Then, the proposed CAD system is covered in section III giving more attention to the description of LBP features and the 3D extension. Next, section IV presents and discusses the experimental results and, finally, a brief summary of the most important conclusions is given in section V.

II. RELATED WORK

In the last decade, a large number of works studying different systems for the computerized diagnosis of AD and MCI have been published, where the discriminative information is provided by at least one of the following neuroimaging techniques: Single-Photon Emission Computed Tomography (SPECT), Magnetic Resonance Imaging (MRI) or PET. The baseline approach uses the Voxel Intensities (VI) of these neuroimages directly as features, and most studies differ essentially in the techniques used to obviate the *curse of dimensionality*. For instance, in [10] previous knowledge about the disease was explored to segment the regions that are typically mostly affected by AD and only the average intensities of such regions of interest were used as features. In [11], two multivariate techniques, Principal Component Analysis (PCA) and Linear Discriminant Analysis (LDA), reduced substantially the dimensionality of the problem through a linear combination of the input features and then a Bayesian framework was used for learning. The best results were obtained using LDA because this technique focuses on the separation of subjects of different classes during the dimensionality reduction stage.

Alongside with the voxel intensities, texture is also a very important property of images, which typically also contains useful information. However, only a few works have previously studied its discriminative power for the diagnosis of AD or MCI, but obtaining promising

results. Histograms of gradient magnitude and orientation and Haar-like features were used in [12] to extract the texture of FDG-PET images in order to distinguish AD patients, MCI patients and normal controls with very interesting performances. On a different work, features extracted from MR Images using co-occurrence matrices were studied [13]. These features were found to be correlated with the score of the Mini Mental State Exam (MMSE) which is a test for cognitive impairments typically used in AD medical diagnosis.

Beyond the two aforementioned texture descriptors, many others are known such as wavelet coefficients or Laws' texture energy measures. The interested reader is referred to [14] and [15] and references therein for a comparative study and detailed descriptions of a large number of descriptors including all of the above mentioned. Herein, we focused on the Local Binary Patterns [16] which, to our knowledge, were only applied once to the CAD of dementia. In [17], the original two-dimensional LBP version was used to extract the texture of white matter lesions on MR Images with the goal of diagnosing patients suffering from different types of dementia, including but not restricted to Alzheimer's disease.

A 3D version of LBPs should however be used in the problem at hand, in order to match the 3D nature of neuroimages. This generalization has revealed to be an hard problem because no simple, straightforward approach is able to extend two essential LBP concepts: uniformity and rotation invariance. Some attempts have previously been made to surpass these difficulties. In [18], an extension denoted "volume LBPs" was proposed and applied to face recognition in 2D time series (video). However, volume LBPs were originally designed for dynamic texture analysis, i.e. with an explicit time variable, and not for full 3D data. In [19], a full 3D and rotation invariant LBP version was proposed, where LBPs were computed in the frequency domain using spherical harmonics which allowed to build a rotation invariant feature. However, a few simplifications were made. In fact, the sphere was not sampled (a continuous spherical function obtained from the gray-levels of neighbor voxels was considered instead), the threshold operation was replaced by a simple subtraction and the uniformity concept proposed was dataset dependent. Some of these simplifications were later removed in [20], but the uniformity concept was ignored.

III. APPROACH

In this section, we describe each component of the system that was built for the CAD of AD and MCI. An overview is presented in Figure 1. The system was decomposed into four independent components. The first one, normalization, aims to produce a more uniform set of images in order to allow for meaningful comparisons between images of different subjects and obtained from different PET scanners. In fact, this preprocessing step was previously conducted by investigators from the Alzheimer’s Disease Neuroimaging Initiative (ADNI). A brief description of the methodology applied for normalization purposes is given in section III-A. The second component, texture extraction, which is the central topic of this paper, was implemented using a 3D version of Local Binary Patterns. Section III-B1 focuses on the fundamental concepts (in two dimensions) behind LBPs and section III-B2 covers in detail the proposed extension. The third module, feature selection, is often explored to tackle the *curse of dimensionality* by reducing the number of features used for training. Correlation coefficients were chosen to quantify the relevance of each individual feature. The fourth and last component, the learning machine, is responsible for the construction of a model from the input data that can be used subsequently to predict the class of new unseen subjects. The Support Vector Machine (SVM) algorithm was chosen as the learning machine. A brief discussion of the last two components is given in section III-C.

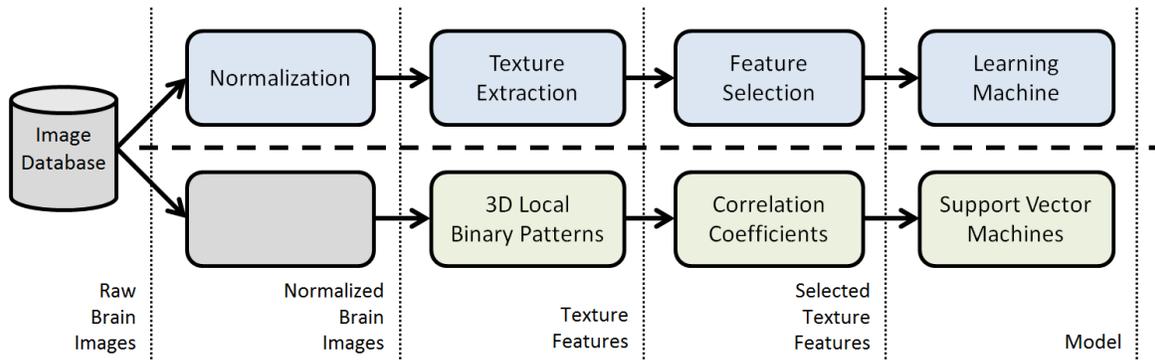


Figure 1: Overview of the CAD system. The top row indicates the purpose of each system component. The bottom row specifies which methods were used to meet them. It should be noticed that the starting database of this work was formed by a set of already normalized PET images. That is the reason why the methods used within the normalization step were omitted in this figure.

A. Dataset and image normalization

The FDG-PET scans used in the current study were retrieved from the ADNI database. Alzheimer’s Disease Neuroimaging Initiative is a large multisite study which, among other goals, has focused on the collection and analysis of different types of neuroimages [21]. In addition to the raw images, a normalized version of all images of this database is available, providing a more consistent starting point for subsequent research. The main goal of this preprocessing step is to eliminate meaningless differences between neuroimages caused, for instance, by different PET scanner models or by anatomical differences in the brain of different subjects. The methodology used for normalization included the following steps: co-registration, orientation alignment and resolution and intensity standardization [22]. The resulting volumes have a dimension of $128 \times 128 \times 60$ voxels, with intensities that span the $[0, 32700]$ interval of integer values.

From all the images available in the database, only a subset of normal controls (NC) with a Clinical Dementia Rating (CDR) of 0, MCI patients with a CDR of 0.5 and AD patients with a CDR of 0.5 or higher were considered in order to limit the number of possible outliers. CDR is a scale used to quantify the stage of dementia, which ranges from 0 to 3, with 0 meaning its absence and 3 the presence of severe symptoms. A brief summary of important clinical and demographic information about each one of the three classes is given in Table I. All statistics refer only to the subset of selected patients.

Table I: Characteristics of each group. Format: Mean (Standard Deviation).

Attributes	AD	MCI	NC
N ^o of subjects	59	59	59
Age	78.3 (6.6)	77.7 (6.9)	77.4 (6.6)
Sex (% of Males)	57.6	67.8	64.4
MMSE	19.6 (5.1)	25.8 (3.0)	29.2 (0.9)
CDR	≥ 0.5	0.5	0

B. Texture extraction

1) *Two-dimensional LBPs*: LBPs [16] were originally designed to differentiate textures in 2D images and are based on a simple element, the Local Binary Pattern, which is constructed in the following way. Consider, first, a set of P points with coordinates $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_P\}$ that

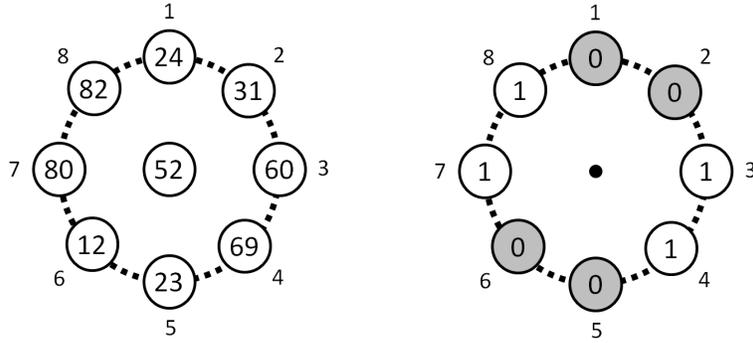


Figure 2: Extraction of a Local Binary Pattern. On the left, the gray-levels of the center and of all neighbor points are illustrated. On the right, the corresponding LBP is shown. The thresholding operation is based on the comparison between the gray-level of each sampling point with the central pixel. The numbers outside the circles represent the indexes to the vector T (Equation (1)).

samples a circumference of radius R , centered at a given pixel with coordinates \mathbf{x}_c . The gray-level of the p -th sampling point will be denoted by V_p and the gray-level of the central pixel by V_c . Notice that the location of each sampling point \mathbf{x}_p is dependent on the center \mathbf{x}_c , and that the gray-level V_p of every sample positioned at non-integer pixel coordinates has to be interpolated. Bilinear interpolation can be applied in these situations. A simple LBP is then computed by thresholding the gray-level of each sampling point with the gray-level of the center, as exemplified in Figure 2. The resulting “bits” can then be joined into a P -dimensional binary vector, T , preserving their angular ordering, i.e. the order of the samples when the circumference is transversed circularly. More precisely, T is defined by

$$T = [H(V_1 - V_c), H(V_2 - V_c), \dots, H(V_P - V_c)]^T, \quad (1)$$

where $H(\cdot)$ is the Heaviside function. In addition, if one thinks of T as a P -bit binary number, it is possible to uniquely identify each LBP by the corresponding (decimal) value. Finally, after computing the LBPs associated with all pixels (by varying the center \mathbf{x}_c), the texture of the image is extracted using the relative frequencies of occurrence of each pattern, i.e. using an histogram.

Despite the simple formulation of this texture descriptor, an obvious problem can be identified. Since the number of distinct patterns grows exponentially with the number of sampling points (2^P to be precise), then the estimations of the true probabilities of occurrence become

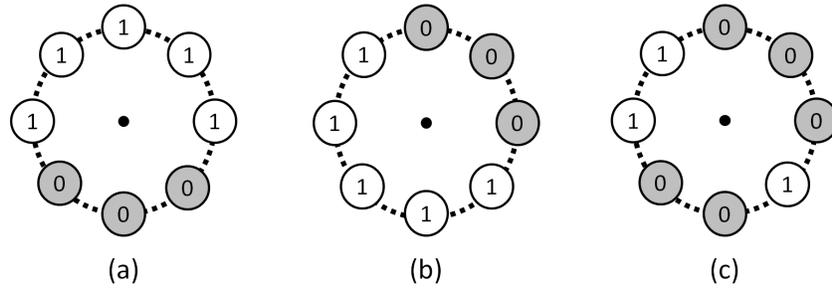


Figure 3: Three examples of LBPs. LBPs (a) and (b) differ only by a rotation and will be merged under the same label, if rotation invariance is considered. LBP (c) can not be aligned with the previous two and therefore will be labeled differently. On the other hand, LBPs (a) and (b) are uniform (presenting two transitions from 1 to 0 or vice versa), while (c) is non-uniform (four transitions).

unreliable for large values of P because not enough samples of each LBP can be extracted from a fixed size image. Consequently, two extensions that reduce the number of differently labeled patterns were proposed: uniformity and rotation invariance.

a) Uniformity: An LBP is said to be uniform (U-LBP) if and only if there are at most two transitions from 1 to 0 or vice versa, when the vector T is transversed circularly [16]. The number of labels is substantially reduced because all non-uniform patterns are merged into a single label. In addition, non-uniform patterns are typically very rare [16] and thus it is unlikely that their probabilities of occurrence contain discriminative information.

b) Rotation invariance: Rotation invariant LBPs (RI-LBPs) do not make any distinction between rotated patterns. Thus, when rotation invariance is used, two LBPs are merged into the same label whenever they can be aligned after an appropriate rotation [16]. Note also that in order to minimize the number of LBP labels, sampling points should be equally spaced over the circumference. In other words, consecutive samples should always be separated by an angular distance of $2\pi/P$.

A few illustrative examples of both uniformity and rotation invariance concepts are given in Figure 3. To conclude the review, it should be stressed that this section focused mainly on the most important concepts associated with LBPs and not on their implementation. For implementation details, the interested reader is referred to the original work [16].

2) Three-dimensional LBPs: FDG-PET scans produce 3D volumes of the brain. Therefore, it is plausible that a 3D extension might improve the descriptor’s discrimination ability.

Simple LBPs can be easily extended to 3D data using a straightforward approach. First,

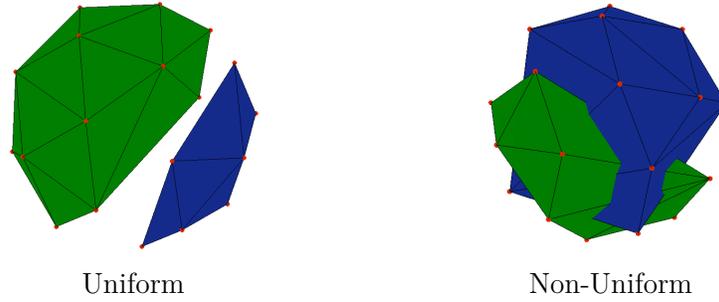


Figure 4: Examples of 3D LBPs. Green hull – \mathcal{H}_0 ; Blue hull – \mathcal{H}_1 . The hulls are disjoint in uniform LBPs and intersect in non-uniform ones.

we sample a sphere of radius R using P points. The notation used for the coordinates and gray-levels of all points will remain the same. Then, after fixing an ordering of the samples, an LBP is encrypted by the binary vector T as in equation (1). However, since in this case there is no natural ordering of the sampling points, we encourage the reader to think of any LBP as the set of all P bits lying on the sphere at the corresponding locations \mathbf{x}_p .

However, a straightforward approach can not extend to three dimensions neither uniformity nor rotation invariance and ,thus, a more elaborated approach is proposed.

a) Uniformity: The main obstacle for the extension of uniformity to 3D data is the definition itself. To solve this problem, we propose a new but equivalent definition: an LBP is said to be uniform if and only if the convex hull \mathcal{H}_0 formed by all neighboring points where $V_p \leq V_c$, and the convex hull \mathcal{H}_1 formed by the remaining ones do not intersect. Two examples are given in Figure 4. In fact, when the new definition is used in two dimensions, the same notion of uniformity is obtained, but now this concept can be easily applied to any higher dimensional space.

In order to efficiently check the uniformity of a given pattern, it should be noted that the hulls \mathcal{H}_0 and \mathcal{H}_1 are disjoint if and only if the sampling points that form them are linearly separable, which is a direct consequence of the Separating Hyperplane Theorem [23]. An overview of a few efficient algorithms that check the linear separability of two sets of points is given in [24].

b) Rotation invariance: First of all, it should be noticed that the equidistant sampling of the sphere, required to reduce significantly the number of rotation invariant LBP labels, is a difficult problem known as Fejes Toth’s. In fact, there is no exact solution for most numbers of sampling points. However, some numerical approximations can be used instead, for instance

the ones proposed in [25].

The principal issue regarding rotation invariance is computational because it is unfeasible to explicitly query against all possible rotations, only to decide whether two patterns can be aligned. We propose the following approach based on spherical harmonics. First, for each LBP, consider the spherical function $f(\theta, \varphi)$ with value one on a small neighborhood over the sphere (with area A) of every point \mathbf{x}_p for which $V_p \geq V_c$ and zero everywhere else, i.e.

$$f(\theta, \varphi) = \begin{cases} H(V_p - V_c) & , \|\mathbf{x} - \mathbf{x}_p\| \leq \varepsilon \quad \forall p \\ 0 & , \text{otherwise} \end{cases} \quad (2)$$

In the previous equation, \mathbf{x} is restricted to the sphere where it can also be represented by the spherical coordinates (θ, φ) . Then, the function $f(\theta, \varphi)$ is decomposed into a linear combination of spherical harmonics

$$f(\theta, \varphi) \approx \sum_{l=0}^{l_M} \sum_{m=-l}^l a_{lm} Y_l^m(\theta, \varphi), \quad (3)$$

where Y_l^m is the spherical harmonic base function of degree l and order m and a_{lm} is the corresponding complex coefficient. Note that the decomposition in equation (3) was truncated with a maximum degree of expansion l_M , but as l_M tends to infinity, the error of reconstruction tends to zero and equation (3) becomes an equality. As an example, the function $f(\theta, \varphi)$ associated with the pattern where all bits are set to 1, and its construction are presented in Figure 5. Finally, the rotation invariant descriptor

$$\text{SH} = \{\|\pi_0(f)\|, \|\pi_1(f)\|, \dots, \|\pi_{l_M}(f)\|\}, \quad (4)$$

proposed in [26] was used to describe the function $f(\theta, \varphi)$, where $\pi_l(f)$ denotes the projection of $f(\theta, \varphi)$ onto the subspace formed by all spherical harmonics of degree l , i.e.

$$\pi_l = \sum_{m=-l}^l a_{lm} Y_l^m(\theta, \varphi). \quad (5)$$

Consequently, provided that the maximum degree of expansion is set high enough, the same SH descriptor is obtained for all patterns that differ only by a rotation, but different ones are computed otherwise. Therefore, it becomes possible to use RI-LBP labels based on this descriptor.

An important implementation detail should now be addressed. Spherical harmonics form an

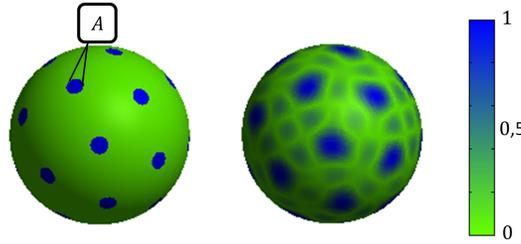


Figure 5: Shape corresponding to the LBP where $V_p \geq V_c$ for all samples p . The exact function is presented on the left and its reconstruction using a small maximum degree of expansion is presented on the right.

orthonormal basis for spherical functions. Using this property, it is possible to express each term of the SH descriptor as a function of the coefficients a_{lm} as follows:

$$\|\pi_l(f)\|^2 = \left\| \sum_{m=-l}^l a_{lm} Y_l^m(\theta, \varphi) \right\|^2 \quad (6)$$

$$= \iint_{\Omega} \left(\sum_{m=-l}^l a_{lm} Y_l^m(\theta, \varphi) \right) \left(\sum_{n=-l}^l a_{ln} Y_l^n(\theta, \varphi) \right)^* d\Omega \quad (7)$$

$$= \sum_{m=-l}^l \sum_{n=-l}^l a_{lm} a_{ln}^* \iint_{\Omega} Y_l^m(\theta, \varphi) Y_l^n(\theta, \varphi)^* d\Omega \quad (8)$$

$$= \sum_{m=-l}^l a_{lm} a_{lm}^* \iint_{\Omega} Y_l^m(\theta, \varphi) Y_l^m(\theta, \varphi)^* d\Omega \quad (\text{Orthogonality}) \quad (9)$$

$$= \sum_{m=-l}^l a_{lm} a_{lm}^* \quad (\text{Normality}) \quad (10)$$

where the * notation stands for the complex conjugate of a number. Additionally, as the area A around each sampling point \mathbf{x}_p tends to zero, the spherical harmonics base functions restricted to that small region becomes approximately constant. Consequently, it is possible to efficiently approximate the coefficients of the expansion by

$$a_{l,m} \equiv \iint_{\Omega} f(\theta, \varphi) Y_l^{m*}(\theta, \varphi) d\Omega \quad (11)$$

$$\approx A \cdot \sum_{p=1}^P H(V_p - V_c) Y_l^{m*}(\theta_p, \varphi_p), \quad (12)$$

where (θ_p, φ_p) is the location of the p -th neighbor in spherical coordinates. When expression (12) is replaced into (10) and subsequently into (4), the factor A will appear in every term of SH and, thus, it can be factored out without harming the discrimination ability or the rotation invariance of the descriptor. Therefore, the actual value of A is not relevant and it can be set

arbitrarily small so that equation (12) holds as an equality in the limit.

On a different topic, the approximations used for the equidistant sampling of the sphere affects the number of patterns that can be merged under the same label because for most sampling schemes there are fewer rotations (if any) that perfectly align all sampling points. In order to enable small displacements of the samples after a rotation and, thus, to reduce the final number of labels, a small difference between the SH descriptors is allowed. Concretely, if one thinks of SH as a vector of dimension $l_M + 1$, then two LBPs described by SH_i and SH_j are considered to be invariant under rotation if

$$\frac{\|SH_i - SH_j\|}{\max\{\|SH_i\|, \|SH_j\|\}} \leq \eta, \quad (13)$$

where η is a parameter that limits the error of misalignment. Additionally, if a given pattern lay within this margin of two distinctly labeled LBPs, then the first is assigned to the group of the closest one. The closeness criterion was defined as in the left-hand side of the previous inequality. This flexibility should not however be allowed for known sampling schemes such as the vertices of a cube for $P = 8$ or the vertices of an icosahedron for $P = 12$.

3) *Texture extraction procedure:* Several issues regarding the application of Local Binary Patterns to neuroimaging data were identified and are discussed in this section.

Typically, rotation invariance is a desired property of texture descriptors. However, all brain images used in this work were previously aligned, which suggests that rotation invariant descriptors might not be used in order to maximize discrimination ability. But, on the other hand, the uncertainty associated with the estimation of the histogram increases if rotation invariance (or uniformity) is not taken into account, which would certainly jeopardize the system's performance. Consequently, we used always both extensions, i.e. U-RI-LBPs.

Different regions of the brain might contain different textures. Therefore, if the probabilities of occurrence of different patterns are computed using all LBP instances extracted from the whole brain image, this spatial information is lost. We limited this phenomenon by splitting the brain volume into a mesh of disjoint cubes of fixed dimension a . Different histograms were then computed inside each cube. Several cube dimensions were tested to allow for the identification of textures at different scales.

Originally, in the 2D version of LBPs, the texture extraction procedure was accelerated

through the use of a look-up table that mapped each one of the 2^P patterns to a uniform and rotation invariant LBP label. Here, we follow a similar approach. However, the construction of this look-up table imposes a computational limit on the number of sampling points, since both the memory required to store the table and the time spent to label all 2^P patterns grows exponentially with P . The memory constrain is easily circumvented because we only need to know the entries of the table associated with uniform patterns and the label of all non-uniforms. As for the timing constrain, an “online” approach would alleviate the problem. In other words, the look-up table can be constructed step-by-step as new patterns appear during the extraction procedure. This approach also has the advantage of not analyzing patterns that are not present in the database. Nevertheless, the values of P tested in this study allowed the “offline” construction of the look-up table.

The standard 2D LBPs were also tested in this study for comparison purposes. They were applied to all axial cuts of the PET image, and the same mesh that was previously described was also used to partition the brain volume into several regions. An histogram was also computed inside each cube.

To summarize, the texture extraction procedure (using either 2D or 3D LBPs) can now be fully stated. First, the look-up table that associates a uniform and rotation invariant LBP label to each pattern is created. Then, for each subject, every brain position is coded with an LBP label and, finally, the probabilities of occurrence of different labels within each cube of the mesh are estimated using an histogram. All entries of all histograms are used as features.

C. Feature selection and learning machine

Many image based CAD systems are characterized by a large disproportion between the number of subjects available for training and the number of features that describe each subject. The feature extraction procedure just presented is no exception. The comparatively small sample size is known to lead to poorer generalization due to a phenomenon known as the *curse of dimensionality* [27], especially if a generative model is used for learning. We tackled this problem in two different ways.

First, we reduced the dimensionality of the feature vector using a feature selection scheme based on the correlation coefficient between each feature and the class label. Thus, from all retrieved texture features, only the ones with highest correlation (in absolute value) with the

class were actually used at the training stage. The number of selected features, N , was left as a parameter to optimize so that the best subset of features (the one that leads to the best performance) can be searched for.

Second, a Support Vector Machine [28], which is a discriminative model, was used for learning purposes. Contrarily to generative models, SVM does not try to learn the probability density functions that generate the data. In fact, this estimation can become unreliable when the number of parameters exceeds the number of training instances. Instead, the SVM algorithm focuses directly on the classification problem at hand by searching for the separation surface that maximizes the margin between subjects of different classes. Consequently, SVM is more robust to the *curse of dimensionality* and it is able to achieve good performances with smaller training sets [29]. Both linear and RBF kernels were tested in the current work.

IV. EXPERIMENTS

In this section, the most important results associated with the texture extraction procedure and the CAD system will be presented. First, the experimental setup will be carefully described in section IV-A. Then, we will validate the proposed extension of the Local Binary Patterns for 3D data in section IV-B and, in section IV-C, the influence of the parameters associated with the new texture descriptor on its discrimination ability will be discussed. Finally, the performance of the CAD system on the diagnosis of both AD and MCI will be presented in section IV-D.

A. Experimental design

The CAD system described in section III will now be used to perform the diagnosis of AD and MCI. In fact, we will test three different approaches for comparison purposes. The first two extract the texture of the image. More concretely, they use the proposed 3D LBP extension and the standard 2D LBP applied to each slice of the image. The last one is not based on texture and uses the voxel intensities of the FDG-PET scan directly. As stated before, this is the baseline approach for image based CAD systems.

Several parameters still need to be set or tuned. When using 3D LBPs, we do not restrict the feature vector to features retrieved from a single resolution. In fact, three values for the radius R of the sphere were allowed, specifically $R \in \{2, 4, 6\}$. Regardless of the radius, the number of

sampling points was always set to 24, which represents the highest number of sampling points for which it was possible to build the look-up table in an acceptable amount of time. The mesh used to partition the brain was also tuned by varying the size, a , of the basic cube in the range $\{9, 13, \dots, 33\}$. All texture features extracted using every parameter setting (R, P, a) were concatenated to form the feature vector that entered the feature selection stage. The parameter η was fixed to 0.05.

As for 2D LBPs, the same meshes were used to partition the brain. However, different values for the pair (R, P) were tested, specifically $(R, P) \in \{(2, 16), (4, 30), (6, 48)\}$. It should be noted that in the 2D case, the construction of the look-up table imposes no computational limitation on the value of P because it is possible to enumerate all uniform patterns without having to analyze non-uniform ones. When VIs were used directly (without texture extraction), only intracranial voxels were considered.

Additionally, the number of features retained in the feature selection stage and the parameters associated with the SVM algorithm were optimized using a grid-search approach conducted within a 10×10 -fold nested cross-validation procedure in order to estimate in an unbiased fashion the performance of the CAD system. N was allowed to assume any value from the set $\{50, 100, 250, 500, 1000, 2500, 5000, 10000, 25000, 50000\}$. As for the SVM parameters, both C (the parameter that controls the cost of misclassification) and γ (the parameter that shapes the RBF kernel) were tuned within a very large range (from 2^{-18} to 2^{18}) using a geometrical progression.

B. Validation of the proposed LBP extension

As mentioned before, both uniformity and rotation invariance were proposed with specific goals. The uniformity concept aims to reduce substantially the number of labels by merging the least frequent patterns (the non-uniform ones) into one label. Rotation invariance reduces even more the number of labels and allows the identification of similar but differently oriented textures. Tables II and III present three statistics that allow the assessment of the proposed extensions of both concepts to 3D data.

The first two statistics, number of uniform LBPs and their combined incidence rate, show that the redefinition of the uniformity concept accomplishes its purpose. First, the large majority of LBPs are indeed non-uniform, reason why the number of U-LBPs is so small when compared

Table II: Statistics of 2D LBPs.

P	Incidence		
	Number of U-LBPs	Rate of U-LBPs ($R=2, R=4, R=6$)	Number of U-RI-LBPs
16	242	(92.1, 79.7, 71.5)%	18
32	1 262	(91.5, 77.9, 67.5)%	34
48	2 258	(90.9, 76.9, 66.8)%	50

Table III: Statistics of 3D LBPs.

P	Incidence		
	Number of U-LBPs	Rate of U-LBPs ($R=2, R=4, R=6$)	Number of U-RI-LBPs
6	46	(96.3, 95.3, 94.9)%	8
8	104	(92.8, 91.8, 91.7)%	11
12	338	(87.7, 84.7, 83.5)%	15
20	1 578	(83.5, 79.5, 77.8)%	30
24	3 412	(82.3, 78.0, 76.2)%	96

to the number of possible patterns (2^P to be exact). For instance, in the 3D case with $P = 24$, there are only 3412 uniform patterns out of 2^{24} possible ones, i.e., 99.98% of all LBPs are non-uniform and, consequently, they are merged into the same label. Nevertheless, the patterns that were identified as uniform are indeed the most frequent ones as can be seen by their combined incidence rate in our database, which represents the relative frequency of occurrence of uniform patterns when considering all patterns extracted from all intracranial voxels of all subjects' PET images. As an example, if one considers the same configuration, 3D LBPs with $P = 24$, it is remarkable that, in the worst scenario presented, 76.2% of all extracted patterns are uniform, especially because the group of uniform patterns represent only 0.02% of all differently shaped ones ($3412/2^{24}$).

The third statistic is the number of distinct LBP labels when both uniformity and rotation invariance are used. In fact, the features resulting from the texture extraction procedure represent the probabilities of occurrence of these labels on different regions of the brain. Thus, since the number of LBP instances extracted from neuroimages and available to perform such estimations is limited (and equal to a^3 , to be specific), the number of U-RI-LBP labels should be small enough so that good estimations can be computed, at most a few hundreds given the

values of a tested in this work. This restriction is met in both 2D and 3D case for all numbers of sampling points. Therefore, the rotation invariance concept also meets its purposes, since it is not only able to group differently oriented patterns, but most importantly because it reduces the number of LBP labels, allowing reliable estimations of their probabilities of occurrence. Finally, notice that the number of U-RI-LBP labels increases considerably when 24 samples are considered in the 3D case. This is a consequence of the approximation used for the sampling configuration of the sphere. In fact, all other settings listed in Table III use known configurations that correspond to the coordinates of the vertices of some regular 3D polyhedron which, in turn, are known to have a large number of rotational symmetries (i.e. number of rotations after which the polyhedron looks exactly the same).

C. Influence of LBP parameters

The influence of the parameters associated with the texture extraction procedure (i.e., a and R) in the diagnosis of AD was studied. A similar analysis was also conducted for the diagnosis of MCI but no relevant differences were found and, therefore, the results were omitted.

Figure 6 shows two series for both parameters a and R . The first series (in red) represents the percentage of features extracted from the database that are associated with each parameter value. The second series (in green) represents the same percentage but instead of considering the set of all extracted features, each percentage is computed considering only the subset of 1% of the features with highest correlation (in absolute value) with the class label.

Regarding the size, a , of each brain partition, it is clear that the selection procedure favors larger regions for the computation of the histograms. Although features extracted from cubes with minimal dimension ($a = 9$) are still the most frequent after selection, the setting $a = 9$ is the only one for which its relative contribution decreases. In fact, from the 10 features with highest correlation with the class label, the settings $a = 13$ and $a = 17$ are responsible for 8 and the setting $a = 9$ only for 1. This fact is consistent with an observation previously made: since the number of instances available for the estimation of the probabilities of occurrence of all LBP labels is small when $a = 9$ ($9^3 = 729$), then the uncertainty associated with those estimations is significant and their discriminative power can only decrease.

On the other hand, no value of radius R of the sphere was systematically preferred or rejected by the selection algorithm, which indicates that this parameter does not influence the

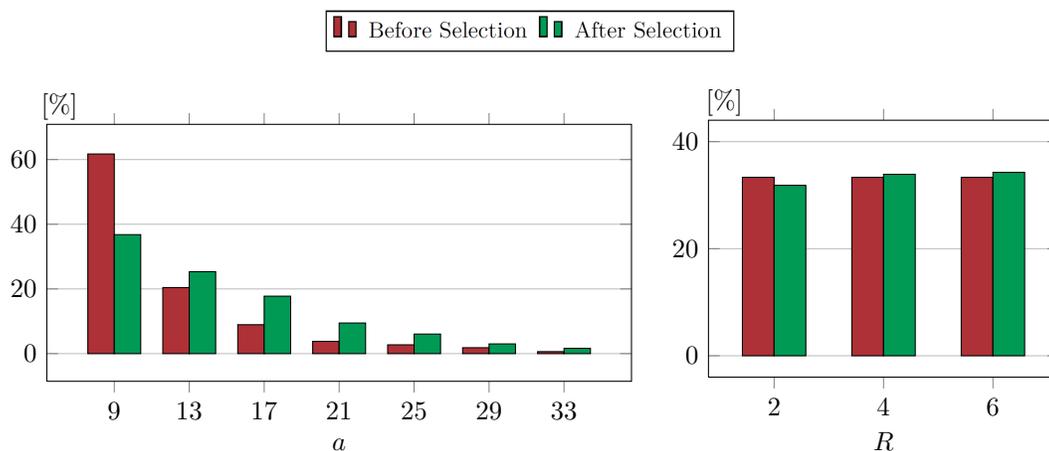


Figure 6: Influence of the parameters a and R on the discrimination ability of a given feature. The series “Before Selection” (in red) shows the fraction of features extracted using each specific value of the parameter, when considering all available ones. The series “After Selection” (in green) shows the same fraction but now, instead of considering all features initially extracted, each fraction was computed within the set of the 1% of features with highest correlation with the class label.

discrimination ability of the texture descriptor much.

D. Classification results

Finally and most importantly, the performances obtained in the diagnosis of AD and MCI were also studied. Table IV presents the accuracy, sensitivity (true positive rate) and specificity (true negative rate) achieved in each setting with the best marks signaled with boldface type.

The novel extension reported good overall results. In fact, it was able to improve significantly the performance attained with the standard 2D LBPs in both diagnostic problems, which validates the initial belief that a 3D version of the texture descriptor would enhance its discrimination power due to the 3D nature of the neuroimaging data. Additionally, it also achieved significantly better results than the typical approach based on VIs in the diagnosis of AD, and similar results (slightly superior) in the diagnosis of MCI. These performances indicate that the texture of FDG-PET images hold discriminative information about the presence of AD, even in its early stages.

Three details should also be stressed. First, the specificity yielded consistently better marks than sensitivity regardless of the type of feature in use and the diagnostic problem at hand, which means that a negative diagnosis, i.e. the decision that a subject is healthy, is more reliable than a positive one. In fact, the 3D LBPs outperformed the other approaches in the AD vs. NC task because it was able to increase significantly the reliability of a positive diagnosis. Second,

Table IV: Classification results using linear and RBF SVM kernels. ACC – Accuracy [%]; SENS – Sensitivity [%]; SPEC – Specificity [%].

		AD vs. NC			MCI vs. NC		
		ACC	SENS	SPEC	ACC	SENS	SPEC
VI	Linear	86.7	83.2	90.2	75.0	70.8	79.2
	RBF	86.0	84.2	87.8	74.5	69.5	79.5
2D-LBP	Linear	88.9	83.7	94.1	71.3	68.8	73.7
	RBF	87.3	86.4	88.1	71.9	69.7	74.2
3D-LBP	Linear	90.5	89.0	91.9	75.6	72.9	78.4
	RBF	89.7	88.0	91.4	73.8	71.5	76.1

the use of the RBF kernel for the SVM algorithm did not improve significantly any of the performance measures, achieving always similar or worse accuracies than the ones attained with a linear kernel, despite its higher computation time requirements. Remember that there is one additional parameter to optimize, γ , when the RBF kernel is being used. Finally, the texture extraction procedure using 3D LBPs was restricted to only 24 sampling points, which can be regarded as a sparse sampling of the sphere, particularly for large radii. However, the experimental results just presented show that this is not a limitation since good generalization abilities can still be achieved.

V. CONCLUSION

In this paper, we proposed a novel generalization of Local Binary Patterns to 3D data. Contrarily to the extensions found in the literature, the proposed approach is able to closely replicate in three dimensions and without any approximation both uniformity and rotation invariance concepts originally proposed for the 2D setting.

We applied the new texture descriptor to the diagnosis of both AD and MCI, yielding interesting results. More precisely, 3D LBPs outperformed its 2D counterpart and the standard approach based on VIs in both classification tasks. These results indicate, on the one hand, that the texture of FDG-PET images contains highly discriminative information about the presence of AD and, on the other hand, that the generalization ability of Local Binary Patterns for the CAD of AD and MCI can be enhanced by the 3D extension.

Finally, it should be stressed that although the proposed texture extraction procedure was applied to FDG-PET images, the same methodology can be applied to other neuroimaging

techniques.

VI. ACKNOWLEDGMENTS

This work was supported by Fundação para a Ciência e Tecnologia (FCT/MCTES) through the ADIAR project (PTDC/SAU-ENB/114606/2009).

Data used in the preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative database. As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report.

REFERENCES

- [1] Z. Khachaturian and T. Radebaugh, *Alzheimer's Disease: Cause(s), Diagnosis, Treatment, and Care*. Taylor & Francis Group, 1996.
- [2] R. Brookmeyer, S. Gray, and C. Kawas, "Projections of Alzheimer's disease in the United States and the public health impact of delaying disease onset," *American Journal of Public Health*, vol. 88, no. 9, pp. 1337–1342, 1998.
- [3] R. Brookmeyer, M. Corrada, F. Curriero, and C. Kawas, "Survival following a diagnosis of Alzheimer disease," *Archives of Neurology*, vol. 59, no. 11, pp. 1764–1767, 2002.
- [4] A. Association, "2012 Alzheimer's disease facts and figures," *Alzheimer's and Dementia: The Journal of the Alzheimer's Association*, vol. 8, no. 2, pp. 131–168, 2012.
- [5] J. Morris, M. Storandt, and J. Miller, "Mild cognitive impairment represents early-stage Alzheimer disease," *Archives of Neurology*, vol. 58, no. 3, pp. 397–405, 2001.
- [6] G. Waldemar, B. Dubois, M. Emre, J. Georges, I. G. McKeith, M. Rossor, P. Scheltens, P. Tariska, and B. Winblad, "Recommendations for the diagnosis and management of Alzheimer's disease and other disorders associated with dementia: EFNS guideline," *European Journal of Neurology*, vol. 14, no. 1, pp. e1–e26, 2007.
- [7] D. H. S. Silverman, "Brain 18F-FDG PET in the diagnosis of neurodegenerative dementias: comparison with perfusion SPECT and with clinical evaluations lacking nuclear imaging," *Journal of Nuclear Medicine*, vol. 45, no. 4, pp. 594–607, 2004.
- [8] K. Herholz, S. Carter, and M. Jones, "Positron emission tomography imaging in dementia," *The British Journal of Radiology*, vol. 80, pp. 160–167, 2007.

- [9] G. Castellano, L. Bonilha, L. Li, and F. Cendes, "Texture analysis of medical images," *Clinical Radiology*, vol. 59, no. 12, pp. 1061–1069, 2004.
- [10] K. R. Gray, R. Wolz, S. Keihaninejad, R. A. Heckemann, P. Aljabar, A. Hammers, and D. Rueckert, "Regional analysis of FDG-PET for use in the classification of Alzheimer's disease," in *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pp. 1082–1085, 2011.
- [11] M. López, J. Ramírez, J. Górriz, D. Salas-Gonzalez, I. Álvarez, F. Segovia, and R. Chaves, "Multivariate approaches for Alzheimer's disease diagnosis using Bayesian classifiers," in *Nuclear Science Symposium Conference Record (NSS/MIC), 2009 IEEE*, pp. 3190–3193, 2009.
- [12] E. Bicacro, M. Silveira, and J. S. Marques, "Alternative feature extraction methods in 3D brain image-based diagnosis of Alzheimer's disease," in *Image Processing (ICIP'12), 2012 IEEE International Conference on*, pp. 134–137, 2012.
- [13] X. Li, H. Xia, Z. Zhou, and L. Tong, "3D texture analysis of hippocampus based on MR images in patients with Alzheimer disease and Mild Cognitive Impairment," in *Biomedical Engineering and Informatics (BMEI), 2010 3rd International Conference on*, vol. 1, pp. 1–4, 2010.
- [14] T. Randen and J. H. Husøy, "Filtering for texture classification: A comparative study," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 21, no. 4, pp. 291–310, 1999.
- [15] L. G. Shapiro and G. C. Stockman, *Computer Vision*, ch. Texture, pp. 235–248. Prentice Hall, 2001.
- [16] T. Ojala, M. Pietikäinen, and T. Mäenpää, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 24, no. 7, pp. 971–987, 2002.
- [17] K. Oppedal, K. Engan, D. Aarsland, M. K. Beyer, O.-B. Tysnes, and T. Eftestøl, "Using local binary pattern to classify dementia in MRI," in *Biomedical Imaging (ISBI'12), 9th IEEE International Symposium on*, pp. 594–597, IEEE, 2012.
- [18] G. Zhao and M. Pietikäinen, "Dynamic texture recognition using local binary patterns with an application to facial expressions," *Pattern Analysis and Machine Intelligence*,

- IEEE Transactions on*, vol. 29, no. 6, pp. 915–928, 2007.
- [19] J. Fehr, “Rotational invariant uniform local binary patterns for full 3D volume texture analysis,” in *Finnish Signal Processing Symposium (FINSIG’07), Proc.*, 2007.
- [20] J. Fehr and H. Burkhardt, “3D rotation invariant local binary patterns,” in *Pattern Recognition (ICPR’08), 19th International Conference on*, pp. 1–4, 2008.
- [21] Laboratory of Neuro Imaging, UCLA, “About ADNI.” <http://adni.loni.ucla.edu/about/>, 2012. [Online: accessed December 1, 2012].
- [22] Laboratory of Neuro Imaging, UCLA, “PET protocols.” <http://adni.loni.ucla.edu/research/protocols/pet-protocols/>, 2012. [Online: accessed December 1, 2012].
- [23] S. Boyd and L. Vandenberghe, *Convex Optimization*, ch. Separating and supporting hyperplanes, pp. 46–50. Cambridge Univ. Press, 2009.
- [24] D. Elizondo, “The linear separability problem: some testing methods,” *Neural Networks, IEEE Transactions on*, vol. 17, no. 2, pp. 330–344, 2006.
- [25] R. H. Hardin, N. J. A. Sloan, and W. D. Smith, “Tables of spherical codes with icosahedral symmetry.” <http://neilsloane.com/icosahedral.codes/>, 2012. [Online: accessed December 1, 2012].
- [26] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz, “Rotation invariant spherical harmonic representation of 3D shape descriptors,” in *Geometry Processing (SGP ’03), Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on*, pp. 156–164, Eurographics Association, 2003.
- [27] G. Hughes, “On the mean accuracy of statistical pattern recognizers,” *Information Theory, IEEE Transactions on*, vol. 14, no. 1, pp. 55–63, 1968.
- [28] C. Cortes and V. Vapnik, “Support-vector networks,” in *Machine Learning*, vol. 20, pp. 273–297, 1995.
- [29] R. P. W. Duin, “Classifiers in almost empty spaces,” *Pattern Recognition, International Conference on*, vol. 2, pp. 1–7, 2000.