EXTENDING LOCAL BINARY PATTERNS TO 3D FOR THE DIAGNOSIS OF ALZHEIMER'S DISEASE

Pedro M. Morgado, Margarida Silveira and Jorge S. Marques

Instituto Superior Técnico, Institute for Systems and Robotics, Lisbon, Portugal

ABSTRACT

Neuroimaging has shown great potential for the computer-aided diagnosis (CAD) of both Alzheimer's disease (AD) and Mild Cognitive Impairment (MCI). However, the texture of such images has been little explored. In this paper, we explore the discriminative power of the Local Binary Patterns (LBPs) texture descriptor to diagnose AD and MCI from 3D brain images. For this purpose, we propose a novel extension of LBPs to full 3D data that, unlike previous approaches, makes no approximations to the underlying concepts of uniformity and rotation invariance. Experimental results obtained using FDG-PET images from the Alzheimer's Disease Neuroimaging Initiative (ADNI) showed that the new feature was able to improve the system's performance when compared to the raw Voxel Intensities (VI) and to the standard 2D LBP version applied to axial cuts of the PET volume.

Index Terms— Computer-Aided Diagnosis, Alzheimer's Disease, Mild Cognitive Impairment, 3D Local Binary Patterns, Positron Emission Tomography, Support Vector Machine

1. INTRODUCTION

Alzheimer's disease is a neurological disorder that mostly affects people over 65 years old and whose incidence rate grows exponentially with age. It is a progressive disease meaning that it worsens over time, affecting memory, cognitive and physical capabilities, and eventually leading to death. Currently, no treatment can cure or stop the progress of AD, but some pharmaceuticals have proven effective in slowing down the advance of symptoms, especially if the disease is detected in its early stages. A syndrome that is proved to be related with the preclinical stage of AD is MCI and, thus, its diagnosis is essential to improve patients' life quality.

Several neuroimaging techniques have demonstrated to hold precious information about the presence of AD, even in its early stages. Consequently, they have been studied as the main input of several fully automated diagnostic systems [1, 2]. In the current paper, we will use the information contained in the texture of the brain image produced by the FDG-PET neuroimaging technique to build a system for the CAD of both AD and MCI.

Texture is an alternative source of information that can be retrieved from an image, and a great variety of tools to extract it are known. Two important examples are histograms of gradient magnitude and orientation (HGMO) and co-occurance matrices which were studied before for the CAD of AD in [1] and [2], respectively. Another approach that recently has been successfully applied to a wide range of different applications, from texture analysis to face recognition, is based on the Local Binary Patterns [3].

LBPs were originally proposed for 2D texture analysis but there have been several attempts to extend them to 3D data. However, all extensions found in the literature introduced some sort of approximation. In [4], "volume LBPs" are proposed, but they deal with dynamic texture analysis of 2D time series and not full 3D data. In [5], not only the neighborhood of a voxel is not thresholded with the central value (the operation is replaced by a simple subtraction) but also the notion of uniformity that is used depends on the dataset. Finally, another scheme for rotation invariant LBPs is presented in [6] but the uniformity concept was ignored. Herein, a novel approach to full three-dimensional, uniform and rotation invariant LBPs is introduced without any approximation of the underlying concepts.

To our knowledge, 3D LBPs were never used for the automated diagnosis of AD. In fact, only one related work was found in the literature where 2D LBPs were used to distinguish between patients suffering from dementia from normal controls [7]. However, LBPs were used in that study for the texture analysis of white matter lesions on MRI, and no effort was done to extend the descriptor to 3D data. Moreover, the study was not restricted to AD patients and also included patients suffering from Lewy Body dementia.

The remainder of the paper is organized as follows: First, in section 2, we review the original 2D LPBs and present in detail the proposed extension to 3D data. Then, the experiments and their results are presented and discussed in section 3. Section 4 concludes the paper.

2. APPROACH

2.1. Two-Dimensional LBPs

Local Binary Patterns [3] were originally proposed for the texture analysis of 2D images. An LBP encodes the texture of the neighborhood of a given pixel $\mathbf{x}_c$ with gray-level V_c using P equidistant neighbors located over a circumference of radius R. The coordinates of the p-th neighbor will be denoted by $\mathbf{x}_p$, and its gray-level by V_p. In addition, neighbors are often located at non-integer coordinates and bilinear interpolation can be used to compute V_p. The encryption is done by thresholding the neighbors with the gray-level of the central pixel, yielding the P-dimensional binary vector

$$T = \left[H\left(V_1 - V_c\right), \ldots, H\left(V_P - V_c\right)\right]^T, \tag{1}$$

where $H(\cdot)$ is the Heaviside function. Vector T can be interpreted as a binary number and, therefore, can be uniquely identified by the corresponding (decimal) value. Then, after labeling all pixels in the image by varying the central pixel $\mathbf{x}_c$, the probability of occurrence of every pattern T is estimated using an histogram. However, the

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 (ADNI) 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.

number of possible patterns grows geometrically with the number of considered neighbors, increasing the uncertainty associated with these estimations. Two extensions that reduce the number of labels were proposed in [3]:

- Uniform LBPs (U-LBPs) – An LBP is said to be uniform if the binary vector T contains at most two transitions from 0 to 1 or vice versa when traversed circularly. All non-uniform patterns are merged into the same label.
- Rotation invariant LBPs (RI-LBPs) – Rotation invariance merges under the same label patterns that can be aligned after an appropriate rotation.

2.2. Three-Dimensional LBPs

A novel approach to full 3D U-RI-LBPs is now proposed. First, consider a 3D equidistant neighbor set lying on a sphere of radius R. In fact, equidistant sampling of the sphere has no exact solution for most numbers of sampling points P (a problem known as Fejes Toth's). However, some numerical approximations are available, as the ones proposed in [8]. Thus, simple LBPs can be encrypted by the binary vector T, as in the 2D case. It is when the concepts of uniformity and rotation invariance are included that obstacles arise.

2.2.1. Uniformity

The original definition of uniformity can not be generalized to higher dimensions and, therefore, a new definition is proposed: an LBP is considered to be uniform if and only if the convex hull $\mathcal{H}_0$ of the sampling points for which $H(V_p - V_c) = 0$ and the convex hull $\mathcal{H}_1$ of the remaining ones do not intersect. Figure 1 illustrates one example of uniform and non-uniform pattern. Note that this definition can be applied directly to the original 2D case, leading to the same notion of uniformity. Now, from the Separating Hyperplane Theorem [9], one can conclude that if $\mathcal{H}_0$ and $\mathcal{H}_1$ are disjoint (do not intersect), then the hulls are linearly separable. Several algorithms that check the linear separability between two sets of points exist in the literature as can be seen in [10], and can be used to determine efficiently the uniformity of any pattern.

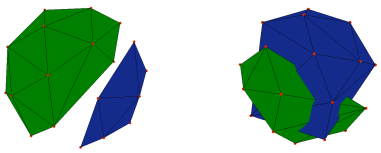

Uniform Non-Uniform

Fig. 1. Examples of LBPs. Green hull – $\mathcal{H}_0$; Blue hull – $\mathcal{H}_1$. The hulls are disjoint in uniform LBPs and intersect in non-uniform ones.

2.2.2. Rotation Invariance

In order to decide whether two patterns can be aligned after a rotation without having to explicitly query against all possible transformations, we first constructed a spherical function $f(\theta, \varphi)$ for each LBP, which is equal to one in a small neighborhood of every point $\mathbf{x}_p$ over the sphere (with area A) for which $H(V_p - V_c) = 1$ and zero everywhere else, i.e.

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

In the previous equation, $\mathbf{x}$ is restricted to the sphere where it can also be represented by the spherical coordinates (θ, φ). Function $f(\theta, \varphi)$ was then decomposed into its spherical harmonics with a maximum degree of expansion l_M, i.e.

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

where Y_l^m is the spherical harmonic of degree l and order m and a_{lm} is the corresponding complex coefficient. Finally, the rotation invariant shape descriptor SH proposed in [11] was used to describe f and, thus, to identify rotation invariant LBPs. SH is given by

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

where π_l represents the projection of f onto the subspace formed by the span of the spherical harmonics with the degree fixed to l, i.e. $\pi_l = \sum_{m=-l}^{l} a_{lm} Y_l^m$. A comprehensive description of spherical harmonics can be found in [12].

Important implementation details should now be addressed. First, each term of SH can be efficiently computed by noting that

$$\left|\left|\pi_l(f)\right|\right|^2 = \left|\left|\textstyle\sum_{m=-l}^{l} a_{lm} Y_l^m\right|\right|^2 \tag{5}$$

$$= \iint_\Omega \Big(\textstyle\sum_{m=-l}^{l} a_{lm} Y_l^m\Big)\Big(\sum_{m=-l}^{l} a_{lm} Y_l^m\Big)^* d\Omega \tag{6}$$

$$= \textstyle\sum_{m=-l}^{l} a_{lm} a_{lm}^*, \tag{7}$$

where the * notation stands for the complex conjugate of a number. Equality (7) is obtained after expanding the product of the sums and using the orthonormal property of spherical harmonics. Additionally, as the area A over the sphere around each point $\mathbf{x}_p$ tends to zero, the function $Y_l^m(\theta, \varphi)$ restricted to that small region becomes approximately constant. Consequently, the harmonic coefficients a_{lm} can be approximated by

$$a_{lm} \equiv \iint_\Omega f(\theta,\varphi) \cdot Y_l^{m*}(\theta,\varphi) d\Omega \tag{8}$$

$$\approx A \cdot \textstyle\sum_{p=1}^{P} H(V_p - V_c) \cdot Y_l^{m*}(\theta_p, \varphi_p), \tag{9}$$

where (θ_p, φ_p) is the location of the p-th neighbor in spherical coordinates. Note also that when (9) is replaced into (7) and then into (4) the factor A is common to all terms in SH. Therefore, its actual value is not relevant and can be set arbitrarily small so that equation (9) holds as an equality in the limit.

On a different note, since equidistant sampling is only approximated for some cardinalities of the neighbor set, which affects rotation invariance, a small difference between the SH descriptors was allowed. More precisely, if one thinks of SH as a vector of dimension $l_M + 1$, the same label was assigned to two LBPs i and j if

$$\frac{||\text{SH}_i - \text{SH}_j||}{\max\{||\text{SH}_i||, ||\text{SH}_j||\}} \leq \eta. \tag{10}$$

Additionally, if a given pattern lies within this margin of two distinctly labeled LBPs, then that pattern is assigned to the group of the closest LBP. The closeness criterion was defined as in the left-hand side of the previous inequality.

2.3. LBPs on Neuroimaging Data

A brain image might not be described by a single texture. Thus, several histograms of LBPs were computed inside cubes that were organized in a mesh that spanned the entire volume. Also, several cube

dimensions, a, were tested allowing for the identification of textures that are present at different scales. On the other hand, the fact that all brain images in our database were aligned suggests that the use of RI-LPBs should not be necessary in order to maximize discrimination ability. However, if we do not consider rotation invariance, the number of possible patterns is very high, and the uncertainty associated with the estimation of the probability of occurrence of each pattern is also considerable, which may jeopardize the system's performance. Consequently, our choice was to use U-RI-LBPs.

The feature extraction procedure can now be fully stated: First, a look-up table that maps each pattern to a U-RI-LBP label is created. In this table, all non-uniform patterns are tagged with the same label, which is different from the ones that identify each group of rotation invariant patterns. Afterwards, each position of the brain image is labeled using the look-up table, and then several histograms are constructed, each one in a different cube of the mesh. Finally, the feature vector is created concatenating all entries of all histograms.

The construction of the look-up table imposes a computational limit on the number of 3D neighbors, since its computation time grows geometrically with P. However, although this limiting step has been performed offline, an online approach would alleviate the problem, i.e. the patterns could be labeled as they appeared in the database which would mean that most non-uniform patterns (absent from the database) would never be analyzed. As for the 2D setting, time complexity is not a concern since we can enumerate all uniform patterns without having to analyze non-uniform ones.

2.4. Feature Selection and Learning Machine

The number of features produced by the extraction procedure described above is extremely high when compared to the small number of subjects available for training. This disproportionality is known to be responsible for the performance degradation of many pattern recognition systems, and it was tackled in two different ways. First, only the most discriminant features were retained, using a ranking selection procedure based on the Pearson correlation coefficient between each feature and the class label, i.e. by selecting only the features with highest correlation with the class. The number of features, N, to select was left as a parameter to optimize. Second, the SVM algorithm was used for learning purposes. SVM [13] is one of the most popular discriminative models for CAD both inside and outside the AD research field since it has good generalization ability and produces good results even in the presence of small training sets [14]. The linear kernel was used in the current work.

3. EXPERIMENTS

3.1. Dataset

Neuroimaging data were retrieved from the ADNI database. The following restrictions were imposed to the Clinical Dementia Rating (CDR) of the subjects considered: 0 for normal controls (NC), 0.5 for MCI patients and 0.5 or higher for AD patients, resulting in a dataset composed by 59 subjects for each group. The retrieved data had already undergone the following preprocessing steps: co-registration, orientation alignment, resolution standardization and intensity normalization, resulting in a $128\times128\times60$ voxel grid with intensities that span the $[0, 32700]$ interval[1]. Table 1 summarizes important clinical and demographic information about each group.

[1]A more detailed description of the pre-processing stage is available at http://adni.loni.ucla.edu/methods/pet-analysis/pre-processing/

Table 1. Characteristics of each group. Format: Mean (Standard Deviation). MMSE stands for Mini Mental State Exam.

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)

3.2. Experimental Design

Both 2D and 3D LBPs were implemented in order to evaluate the performance gain achieved by the new scheme. 2D LBPs were applied to axial cuts of the PET image. A multiresolution texture analysis of the brain was conducted by considering several values for the sphere radius R and for the number of samples P. Specifically, the combinations $(R, P) \in \{(2, 16), (4, 30), (6, 48)\}$ were considered for 2D-LBPs and for 3D-LBPs $(R, P) \in \{(2, 24), (4, 24), (6, 24)\}$. P was chosen to always be as high as possible in the 3D case, i.e. 24 for all tested radii, because, otherwise, the neighbors would be sparsely distributed over the sphere. In addition, the mesh from which different histograms were extracted was also tuned by varying the parameter a in the set $\{9, 13, \ldots, 33\}$. All features extracted from all parameter settings were concatenated to form the feature vector which, after the feature selection stage, was used to train an SVM. Besides 2D and 3D LBPs, the common approach that uses VIs as features was also evaluated for comparison purposes.

Both performance assessment and model selection were performed using a 10×10-fold nested cross-validation procedure, which guarantees an unbiased estimate of the system's performance and is also able to tune N and the SVM parameter C that controls the cost of misclassification. The results were averaged after 10 runs. The parameter N was allowed to assume values within the range 50 to 50000 and C within the range 2^{-18} to 2^{18}. A geometrical progression was used to search these ranges. The parameter η was studied and fixed to 0.05.

3.3. Results

Table 2 presents three statistics for varying numbers of sampling points that characterize the 2D and 3D LBP schemes. Specifically, the number of different uniform patterns, their incidence rate in the database (i.e. the percentage of occurrence of uniform patterns when considering all LBPs extracted from all intracranial voxels of all PET images), and the number of different U-RI-LBP labels.

As mentioned before, the motivation for the uniformity concept is related to the fact that, despite representing only a small fraction of the number of possible patterns $\left(2^P\right)$, uniform LBPs are characteristic of textured images. Table 2 confirms this fact for the 2D case where, for instance, in the setting $(R, P) = (2, 16)$, only 0.37% of patterns (242 out of 2^{16}) are uniform and yet they account for 92.1% of the database. On the other hand, the number of U-RI-LBP labels is small enough and, therefore, reasonable estimations of their probabilities of occurrence can be computed using histograms. Note that each histogram is built inside a cube with a minimum dimension of 9 (maximum dimension of 33), which means than each histogram uses at least 729 (at most 35937) LBP instances. Regarding the 3D extension, Table 2 also shows that the new generalizations of both uniformity and rotation invariance also achieve the goals for which they were proposed. 3D U-LBPs still represent a small fraction of the total number patterns, and still account for the vast majority that can be extracted from the database. Rotation invariance is also able

Table 2. Statistics of 2D and 3D LBPs.

	P	Number of U-LBPs	Incidence Rate of U-LBPs (R=2,R=4,R=6)	Number of U-RI-LBPs
2D	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
3D	8	104	(92.8, 91.8, 91.7)%	11
	12	338	(87.7, 84.7, 83.5)%	15
	24	3 412	(82.3, 78.0, 76.2)%	96

to reduce the number of 3D U-RI-LBP labels to small enough values.

The performances obtained with both 2D and 3D LBPs, as well as with VI features, in the AD vs. NC and MCI vs. NC classification problems are presented in Table 3. It can be seen that the novel 3D scheme performed better than its 2D counterpart in both problems and achieved significantly better results than VI features in the diagnosis of AD and similar results in the diagnosis of MCI, which indicates that the new procedure is able to extract discriminative information about the disease. The improvements were attained especially on the positive side of the diagnostic problem (higher sensitivity). In fact, in both problems, there was always one other method that registered a slightly better specificity than the proposed one. However, it can also be noticed that specificity yielded consistently better results than sensitivity in both classification tasks, which is a desirable property for a first-pass clinical procedure as this one.

Finally, the impact of the radius R associated with each pattern on the system's performance can be seen in Table 4, where each row was obtained by restricting the available LBP features to a predefined radius. The results indicate that the most discriminative features for both classification problems are extracted using a radius of four voxels. In fact, the performance attained in the diagnosis of AD with $R = 4$ was even better than with the multi-resolution approach, and only slightly worst in the diagnosis of MCI (see Table 3), suggesting that the inclusion of additional features with less discriminative power can degrade the system's performance.

4. CONCLUSION

This paper studied the use of texture for the computerized diagnosis of both AD and MCI. Specifically, we proposed a new approach to the extension of Local Binary Patterns to three dimensions. This extension differentiates itself from others found in the literature by not introducing any approximation to the original concepts. In fact, both uniformity and rotation invariance were successfully generalized.

In addition, 3D-LBPs proved to hold discriminative information for the diagnosis of AD, even in its early stages. The new texture extraction procedure achieved the best accuracy in both diagnostic problems, outperforming its 2D counterpart and the common approach based on VI features. These results indicate that not only the texture of FDG-PET images is an important source of information for the problem at hand, but also that the proposed 3D extension is able to enhance the generalization ability of the CAD system.

Finally, it should be stressed that although this work has focused on FDG-PET images, the same methodology can also be applied to other neuroimaging modalities.

Table 3. Classification results using a linear SVM kernel. Format: (Accuracy, Sensitivity, Specificity).

	AD vs. NC	MCI vs. NC
VI	(86.7, 83.2, 90.2)%	(75.0, 70.8, **79.2**)%
2D-LBP	(88.9, 83.7, **94.1**)%	(71.3, 68.8, 73.7)%
3D-LBP	(**90.5**, **89.0**, 91.9)%	(**75.6**, **72.9**, 78.4)%

Table 4. Classification results using 3D-LBPs restricted to patterns with constant radius. Format: (Accuracy, Sensitivity, Specificity).

	AD vs. NC	MCI vs. NC
$R = 2$	(88.0, 86.9, 89.0)%	(69.5, 64.2, 74.7)%
$R = 4$	(**91.2**, **89.8**, **92.5**)%	(**75.3**, **71.7**, 79.0)%
$R = 6$	(89.2, 85.9, 92.4)%	(73.2, 65.9, **80.5**)%

5. REFERENCES

[1] E. Bicacro, M. Silveira, and J. S. Marques, "Alternative feature extraction methods in 3D brain image-based diagnosis of Alzheimer's disease," *Image Processing (ICIP'12), 2012 IEEE International Conference on*, pp. 134–137, 2012.

[2] Hong Xia, Longzheng Tong, Xiaoxia Zhou, Jing Zhang, Zhen Zhou, and Weifang Liu, "Texture analysis and volumetry of Hippocampus and medial temporal lobe in patients with Alzheimer's disease," *Biomedical Engineering and Biotechnology (iCBEB), 2012 International Conference on*, pp. 905–908, may 2012.

[3] 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.

[4] 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.

[5] J. Fehr, "Rotational invariant uniform local binary patterns for full 3D volume texture analysis," *Finnish Signal Processing Symposium, Proc.*, 2007.

[6] J. Fehr and H. Burkhardt, "3D rotation invariant local binary patterns," *Pattern Recognition, 19th International Conference on*, pp. 1–4, 2008.

[7] Ketil Oppedal, Kjersti Engan, Dag Aarsland, Mona K. Beyer, Ole-Bjorn Tysnes, and Trygve Eftestøl, "Using local binary pattern to classify dementia in MRI," *Biomedical Imaging (ISBI'12), 9th IEEE International Symposium on*, pp. 594–597, 2012.

[8] 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].

[9] Stephan Boyd and Lieven Vandenberghe, *Convex Optimization*, chapter Separating and supporting hyperplanes, pp. 46–50, Cambridge Univ. Press, 2004.

[10] D. Elizondo, "The linear separability problem: some testing methods," *Neural Networks, IEEE Transactions on*, vol. 17, no. 2, pp. 330–344, 2006.

[11] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz, "Rotation invariant spherical harmonic representation of 3D shape descriptors," *Geometry Processing, Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on*, pp. 156–164, 2003.

[12] C. Müller, *Spherical harmonics: Lecture notes in mathematics*, vol. 17, Springer, 1966.

[13] Corinna Cortes and Vladimir Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, pp. 273–297, 1995.

[14] Robert P. W. Duin, "Classifiers in almost empty spaces," *Pattern Recognition, International Conference on*, vol. 2, pp. 1–7, 2000.