

Directional Invariance of Co-occurrence Matrices within the Liver

Carl Philips Daniel Li
DePaul Johns Hopkins
University University

Cphilips@students.DePaul.edu Daniel.Li@JHU.edu

Daniela Raicu Jacob Furst
DePaul DePaul
University University

DStan@CTI.DePaul.edu JFurst@CS.DePaul.edu

Abstract

Co-occurrence matrices are one of three texture algorithms commonly used on Computed Tomography (CT) images. In this paper we analyze the directional invariance of Co-occurrence matrices for the purpose of reducing their runtime by reducing the number of directions analyzed without negatively affecting the quality of the texture data extracted.

1. Introduction

Organs can be classified using a variety of methods including spatial location [1], shape descriptors [2], and pixel/region texture models [3], [4]. Geographic and shape-based classifiers can be unreliable; exact organ shape and location fluctuate slightly between patients, and organs are often in close proximity to each other, making it difficult to determine where one organ ends and another begins. In addition to avoiding these issues, pixel-based texture classifiers can not only classify organs pixel by pixel but can also flag possible tumors and abnormalities.

According to Tuceryan and Jain [5] there are four families of texture extraction: statistical, transform/signal processing, structural/geographic, and model-based. Each of these families uses a different approach to extracting texture features. Statistical textures focus on probabilities, counts of specific pixel pair occurrences, etc. The transform textures transform an image in such a way as to emphasize patterns or structures before extracting features. The structural texture approach is actually a subset of the geographic family which analyses image primitives to extract features while model-based textures analyze how an image compares to models that predict pixel values.

In this paper we examine the directional redundancy of Co-occurrence, a statistical texture extraction algorithm. To do this we first use a 3D algorithm to gather data for each direction. Then we compare each direction, identifying which directions produce identical information. We expect that this will allow us to analyze a minimal number of directions in the future while still extracting all relevant data. To compare algorithms we used region-based classification, which is similar to pixel-level classification but is not as computationally expensive. The regions were chosen as small cubes

within the Computed Tomography (CT) volume.

The rest of this paper is organized as follows: background and similar work are discussed in section 2, our methodology is explained in section 3, the results of our experiment are shown in section 4, our conclusions are discussed in section 5, and future work is described in section 6.

2. Background

Feature comparisons have been around for a long time [6] however, it is not a highly researched field with respect to co-occurrence and directionality since for each environment and methodology (e.g. CT scans of soft tissue, radar imagery of ice, etc) there is a finite answer that barring significant technological advances, does not change. Barber and LeDrew [7] found that when using co-occurrence matrices, one could use a directional subset while maintaining high classification scores. However, Barber and LeDrew analyzed synthetic aperture radar (SAR) images of sea ice and used only three directions (0°, 45°, 90°). While Barber and LeDrew's findings are relevant, we are interested in the directional redundancy of co-occurrence matrices using thirteen directions when applied to the 3D human CT data. The authors of this paper introduced the main algorithm for 3D co-occurrence matrices and follow up work consisted only in minor modifications of this algorithm. The base approach to 3D co-occurrence matrices was expanded upon by Kovalev et al. through use of gradient vectors [9]. Furthermore, comparisons of 2D and 3D algorithms used metrics that are difficult to apply to non segmentation applications [8].

In this paper we will use two different families of approaches to find the optimal directions for classification. Those families are feature selection and, feature comparison. Guyon, and Elisseeff compared two types of feature selection [10]; the first was feature ranking of which principle component analysis (PCA) is an example. The other type of feature selection is feature subset selection, of which there are three core methods: Backwards Elimination, Forward Selection, and Step-wise Regression.

As both [10] and [11] state, backwards elimination starts with a set S of all features. It then slowly removes each feature that does not significantly contribute to the accuracy. Once completed one is left with only the

features that actively influence the results. Forward selection is the mirrored counterpart in that it starts with an empty set and slowly adds features. Step-wise Regression is a combination of the two: during each iteration a feature may be add to or removed from the set.

An example of an approach for feature comparison would be the least significant different (LSD) tests. LSD tests compare each feature to all other features, statistically identifying which are identical and which are significantly different.

3. Methodology

3.1 Data

To perform our experiments we first had to create 3D regions from each patient's CT slices. From these 3D regions we extracted texture data using the co-occurrence model. This texture data was then used to classify each 3D region as liver or non-liver using a decision tree classifier. These experiments were run on Pentium 4 computers with 3.19GHz CPU's and 1 gigabyte of RAM each.

For this paper we were given explicit permission to use the 20 patients provided for the MICCAI (Medical Image Computing and Computer Assisted Intervention) 2007 Grand Challenge and later maintained for the Segmentation of the Liver Competition 2007 [12]. Patients were given a contrast agent and scanned using a variety of scanners from different manufacturers using different numbers of detector rows (4, 16, and 64 rows). Due to the variety of manufacturers and number of detector rows; intra-slice spacing varies between .55mm and .8mm with an inter-slice distance of between 1mm and 3mm. There was, however, no overlap between slices. The ground truth was determined by radiological experts and includes all internal tumors and blood vessels.

Using liver and non-liver masks derived from the ground truth segmentation, we extracted as many 20^3 pixel cubes from the original CT volume as possible. We observed that the number of non-liver cubes substantially outnumbered the number of liver cubes so non-liver cubes were randomly selected from each corresponding patient until we had an equal number of liver and non-liver cubes for each patient. The remaining non-liver cubes were set aside and effectively discarded.

As a pre-processing step we separately normalized all voxels within each cube by their mean and standard deviation (z-score normalization) in an attempt to

overcome errors introduced by scanner variation.

3.2 Co-occurrence matrices

Co-occurrence matrices operate by finding repetitive pixel/voxel intensity patterns. This is done by tallying the number of times two different pixel intensities are geographically separated by a set number of pixels in a specific direction. These tally results are saved in a matrix with the results for other pixel intensity pairs separated by the same distance and direction. Usually this analysis is performed on multiple directions and several different distances.

These matrices are often quite sparse so many researchers will rescale the range of grayscale pixel values to a lower number. This not only prevents matrix population sparsity but also reduces noise, something that statistical texture extraction methods are sensitive to. For this paper we rescaled the range of grey levels down to 16. This decision is supported by [9] in which it was empirically determined that 4-16 gray levels are optimal.

Directional post-processing may be performed by averaging [13] or adding the directional data together [9], leaving you with one matrix for each distance. This approach aims to create a pseudo-invariance to direction. However, as we are trying to analyze the various directions, we chose to keep all of the matrices.

Co-occurrence matrices are also commonly used in 2D algorithms, which only use four directions. As CT technology advanced, 3D co-occurrence matrices became an option. Figure 1 below illustrates the directional difference between 2D and 3D co-occurrence algorithms. It should be noted that for ease of viewing the image has been rotated 90° to the left.

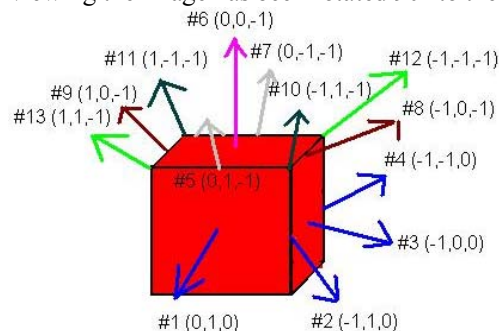


Figure 1. All thirteen directions with corresponding offsets.

Directions 1-4 are the four directions commonly used in 2D co-occurrence matrices while directions 5-13 are the additional nine directions used in 3D co-occurrence matrices. The pixel offset and degree equivalent is shown below in Table 1.

For this paper we used four voxel distances; 1 voxel (neighboring voxels), 2, 4, and 8 voxels. We used a non-linear sequence of distance to span distances from “near” to “far” with a minimal number of distances, which, in turn, increases the processing speed.

Once all the Co-occurrence matrices are created one needs to extract features from the matrices. We used four Haralick features [14]. All of the features used are listed below in Table 2.

Table 1. All thirteen directions

	Offset	Location by degree
1	0,1,0	0°,0° (off the horizontal)
2	-1,1,0	45°,0°
3	-1,0,0	90°,0°
4	-1,-1,0	135°,0°
5	0,1,-1	0°,45°
6	0,0,-1	null,90° (vertical line)
7	0,-1,-1	0°,135°
8	-1,0,-1	90°,45°
9	1,0,-1	90°,135°
10	-1,1,-1	45°,45°
11	1,-1,-1	45°,135°
12	-1,-1,-1	135°,45°
13	1,1,-1	135°,135°

Table 2. Co-occurrence features

	Feature	Description
1	Energy	Measures the occurrence of repeated pairs within an image
2	Correlation	Measures correlation of intensity between pixel pairs
3	Contrast	Measures local contrast
4	Homogeneity	Measures the "purity" of the image

3.3 Principle component analysis (PCA)

The first analysis performed upon the data extracted in section 3.2 was PCA. We organized the data such that each column represented a direction and then partitioned the data by distance, performing PCA on each distance set separately. Finally, we recombined the data and performed PCA for the fifth time. A sample of our results is shown in Table 3 in section 4.

While this test was not imperative, feature ranking tests have the potential to immediately identify important features as it is the case of this paper where we want to identify the most important directions.

3.4 Least significant difference

With the data extracted in section 3.2 we organized it much like in Section 3.3 where each column is associated with a direction. Next the data was

partitioned by feature and distance which separated it into 16 parts. We then performed LSD tests upon these subsets. We then reorganized the data, separating it by distance only and performed LSD tests on these four subsets. The result of one of these tests is shown in Table 4 in Section 4. From these results we were able to group the directions by similarity, for example directions 10-13 are often statistically identical.

In theory, using one direction from each group would provide sufficient data to correctly classify the liver and non-liver cubes.

3.5 Backwards elimination

We chose to use backwards elimination as it is simple and accurate and with only thirteen directions, algorithmic speed was not an issue. We used the classification accuracy (True Positive + True Negative)/(All) for the evaluation stage.. When applied to the data set the direction set (1, 4, 6) (as denoted in Table 1) was selected as optimum.

3.6 Classification

We assigned thirteen patients to the training set and the remaining seven to the testing set. We then used decision trees to classify our results. The particular tree we used was the Classification & Regression (C&RT) tree [15]. There is one parameter for this type of tree; the minimum number of cases required for a nodule to split. We experimented with twelve different parameter values (2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048, and 4096) and selected the one with the best test set accuracy.

4. Results

4.1 Principle component analysis

In Table 3 below, we show the eigenvectors produced by PCA when analyzing all distances. The first column identifies direction and the bottom row identifies how much of the statistical variance of the whole data set is contained within that component. Components 6-13 are not shown as they contain less than 1% of the total variance each.

Table 3. PCA results when analyzing all distances together.

	Comp 1	Comp 2	Comp 3	Comp 4	Comp 5
#1	0.2179	0.1976	0.1017	-0.2599	0.2128
#2	0.2777	-0.4476	0.1105	-0.032	0.2071
#3	0.2487	-0.1994	-0.349	-0.0788	0.2202
#4	0.2727	0.2508	-0.346	0.2302	0.2837
#5	0.2815	0.2604	0.5238	0.0352	-0.0481
#6	0.2406	0.2561	0.1583	-0.4067	-0.3006
#7	0.2615	0.1939	-0.083	-0.4387	-0.2751
#8	0.2938	-0.1992	-0.189	0.1495	-0.4224
#9	0.2851	-0.0071	-0.047	-0.0914	0.2709
#10	0.3153	-0.3506	0.414	0.4068	-0.2464
#11	0.2882	-0.3936	-0.014	-0.359	0.1984
#12	0.2992	0.1436	-0.45	0.2372	-0.3842
#13	0.3067	0.3857	0.1347	0.3631	0.3394
	92.02%	2.09%	1.57%	1.42%	1.09%

Considering that component 1 contains 92.02% of the variance, that is the primary component of interest, and in the above table we see that while most directions contribute similar amounts of data to component 1, directions 1, 3, and 6 contribute the least.

When analyzing the distances separately we found that the amount of data contained within the first component decreases as the distance increases, starting at 97.8% at a distance of 1 voxel and dropping to 88.6% when the distance increases to 8 voxels. However, component 2 is only 3% at that point, still insignificant. It should also be noted that directions 1, 3, and 6 consistently contribute relatively low amounts of variance to component 1 throughout the distances.

4.2 Least significant difference (LSD)

As stated in Section 3.4 we performed 20 LSD tests. In Table 4 we show the results of the LSD test performed upon all features at a distance of 4 voxels.

Table 4. LSD on all features at a distance =4

#1	0												
#2	1	0											
#3	0	0	0										
#4	1	0	0	0									
#5	1	0	1	0	0								
#6	1	0	0	0	1	0							
#7	1	0	1	0	0	0	0						
#8	1	0	1	0	0	1	0	0					
#9	1	0	1	0	0	1	0	0	0				
#10	1	1	1	1	0	1	1	0	0	0			
#11	1	1	1	0	0	1	0	0	0	0	0		
#12	1	1	1	1	0	1	0	0	0	0	0	0	
#13	1	1	1	1	0	1	1	0	0	0	0	0	0
Direction	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	#11	#12	#13

In the above table '1' indicates that the two directions are significantly different while '0' indicates that they are identical. From this table, we can also organize the directions into two groups, (1, 2, 3, 4, 6) and (2, 4, 5, 7-13); note that the two groups do overlap. We then select the primary direction from each group, (3, 8) for table 4, and add it to the list of direction sets to test though classification. A primary direction is defined as the direction within each group that is identical to all group members. In the above table direction one is not a primary direction as it is not identical to any of its group members, save direction three.

From analyzing all 20 tables we found that as the distance increases, so does the variability; meaning that at greater distances, direction is less important. We also found that the higher directions (10-13) tended to be identical, which is interesting as those are the 3D counterparts of directions 2 and 4. While this initially seems to conflict with our PCA results as it seems intuitive to assume that as more directions become statistically identical the directions would become more correlated and thus component 1 would retain more information. However, further analysis of the PCA results showed that as distance increased the coefficient mean for component 1 remained steady. The standard deviation did change; when the distance is one voxel the standard deviation is .064 however it drops to .03 by the time the distance is eight voxels. So while the amount of information retained within component 1 may be reduced as distance increases, the directions contribute more equally.

After we have compiled direction sets to test, we tallied how many times each direction had appeared in that list. This is seen below in Table 5. From this list we found that directions 1, 3, 7, and 12 were represented the most so we added (1, 3, 12) to the list of direction sets to be tested, as (1, 3, 7) was already in the list.

Table 5. Tally of each direction and how many times it was represented.

Direction	Count
1	7
2	2
3	4
4	2
5	1
6	1
7	3
8	1
9	1
10	0
11	2
12	3
13	0

4.3 Classification Results

As stated in section 3.6, classification was performed to determine which direction sets are the best. As a default we first performed the classification on each direction independently, only the 2D directions (1, 2, 3, 4), and all 13 directions. Table 6 summarizes all our accuracy results for all possible combinations. We bolded and italicized the highest accuracy and bolded all accuracies that are statistically identical to it.

From this table we see that many of the direction sets used are statistically identical to (1, 4, 6), the best direction set tested. We also see that when we split (1,4,6) into pairs, only (1,6) is statistically identical. We also see that there are two other statistically identical bi-directional direction sets, and no single direction is statistically identical to (1, 4, 6). This indicates that the optimal direction set has only two directions.

5. Conclusions

In this paper we hypothesized that there is a certain amount of directional variability within the human liver when using co-occurrence matrices, and that this variance would allow one to reduce the number of directions analyzed without reducing the classification accuracy. Our results showed us that one can reduce the number of directions needed from the original four used in 2D co-occurrence matrices down to two without a reduction in classification accuracy. Thus one can significantly reduce the co-occurrence matrices runtime; we found that when analyzing 136 liver cubes, the average runtime for 2D Co-Occurrence matrices (Directions 1,2,3,4) was reduced by 40%, a significant reduction in time.

Table 6. Each direction set tested with its highest accuracy score

Directions	Accuracy
1,2,3,4	0.88177
1,2,3,4,5,6,7,8,9,10,11,12,13	0.88604
1	0.87085
2	0.86657
3	0.87844
4	0.84663
5	0.85138
6	0.85518
7	0.84853
8	0.85043
9	0.85423
10	0.8604
11	0.8528
12	0.8509
13	0.85708
1,6	0.89174
1,2,3,5,6,10	0.88082
1,2,3	0.89221
1,2,6	0.89221
1,2,3,6	0.88604
1,2,5,6,10	0.88841
2,3,5,6,10	0.87939
1,2,4,6,7,10	0.88984
1,2,3,5,6,7,10	0.88224
1,2,5,6	0.88794
2,3,10	0.87987
1,2	0.86515
2,5,12	0.8699
1,6,9	0.89364
1,3,7,12	0.87464
1,4,7,12	0.86467
1,4,11	0.8661
1,3,7	0.87464
1,2,3	0.89079
3,8	0.88604
1,11	0.868
1,3,12	0.89079
3,6	0.89031
1,4,6	0.89459
1,4	0.87037
4,6	0.87417

6. Future Work

The patients we received from the MICCAI Grand

Challenge (2007) are in no way uniform as the voxel sizes vary between patients; this problem could be bypassed by extracting data based on actual volume (in millimeters) as opposed to extracting a set number of voxels. This problem is slowly becoming moot as more and more accurate CT scanners are developed, producing more isometric voxels.

Our results do raise a question however; if co-occurrence matrices have some directional invariance within the human liver, what about other texture algorithms such as Markov Random Fields? We are also interested to explore if these findings apply to all soft tissue within the human body or just liver; two questions we hope to answer in the future.

7. References

- [1] H. Park, P. Bland, C. Meyer, "Construction of an abdominal probabilistic atlas and its application in segmentation," *IEEE Transactions on Medical Imaging*, Vol. 22, No. 4, April 2003, pp. 483-492.
- [2] O. Badaway, M. Kamel, "Shape-based image retrieval applied to trademark images," *International Journal of Image and Graphics*, Vol. 2, No.3 2002, pp. 375-393.
- [3] P. Howarth, S. Ruger, "Robust texture features for still-image retrieval," *IEEE Proceedings on Vision, Image and Signal*, 152(6), 2005, pp868-874
- [4] M. Varma, A. Zisserman, "Texture classification: are filter banks necessary?" *IEEE Proceedings on Computer Vision and Pattern Recognition*, Vol. 2, Madison, Wi, 2003, pp. 691-698.
- [5] Tuceryan M., and Jain A.K., "Texture analysis," *The handbook of pattern recognition and computer vision*, Chen C.H., Pau L.F., and Wang P.S.P, Editors, World Scientific Publishing Co. (1993), pp. 235-276.
- [6] L. Thurstone "Multivariate data analysis". *Psychological Review*, Vol. 38 1931, pp. 406-427.
- [7] D. G. Barber, and E. F. LeDrew, "Science issues relating to marine aspects of the cryosphere: implications for remote sensing." *Canadian Journal of Remote Sensing*, Vol. 18, No. 1, 1991 pp. 46-55.
- [8] A. S. Kurani, D. H. Xu, J. Furst, D. S. Raicu, "Co-occurrence matrices for volumetric data," *The 7th IASTED International Conference on Computer Graphics and Imaging – CGIM 2004*, Kauai, Hawaii, USA.
- [9] V. A. Kovalev, F. Kruggel, H. J. Gertz, D. Y. Cramon, "Three-dimensional texture analysis of mri brain datasets," *IEEE Transactions on Medical Imaging*. Vol. 20, No. 5. 2001.
- [10] I. Guyon, A. Elisseeff, "An introduction to variable and feature selection," *Journal of Machine Learning Research* Vol. 3, 2003, pp. 1157-1182.
- [11] D. Stracuzzi, P. Utgoff, "Randomized variable elimination," *Journal of Machine Learning Research* Vol.5, 2004, pp. 1331-1362.
- [12] B. van Ginneken, T. Heimann, and M. Styner, "3D segmentation in the clinic: A grand challenge", In: T. Heimann, M. Styner, B. van Ginneken (Eds.): 3D segmentation in the clinic: a grand challenge, pp. 7-15, 2007.
- [13] Q. A. Holmes, D.R. Nuesch, and R. A. Shuchman, "Textural analysis and real-time classification of sea-ice types using digital SAR data," *IEEE Trans. Geosci. Remote Sensing*, vol. GE-22, pp. 113-120, Mar. 1984.
- [14] R. Haralick, K. Shanmugam, I Dinstein, "Textural features for image classification." *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC-3, No. 6, November, 1973, pp. 610-621.
- [15] Mathworks, *Matlab R2007a Statistics toolbox* (<http://www.mathworks.com/products/statistics/>), February 1, 2008.