

Evaluation of Binning Strategies for Tissue Classification in Computed Tomography Images

Stefanie Handrick^a, Bahare Naimipour^b, Daniela Raicu^c, Jacob Furst^c

^aArizona State University, Tempe, AZ, USA, 85287

^bUniversity of Illinois, Chicago, IL, USA, 60607

^cIntelligent Multimedia Processing Laboratory

School of Computer Science, Telecommunications, and Information Systems

DePaul University, Chicago, IL, USA, 60604

ABSTRACT[♦]

Binning strategies have been used in much research work for image compression, feature extraction, classification, segmentation and other tasks, but rarely is there any rigorous investigation into which binning strategy is the best. Binning becomes a "hidden parameter" of the research method. This work rigorously investigates the results of three different binning strategies, *linear binning*, *clipped binning*, and *nonlinear binning*, for co-occurrence texture-based classification of the backbone, liver, heart, renal, and splenic parenchyma in high-resolution DICOM Computed Tomography (CT) images of the human chest and abdomen. *Linear binning* divides the gray-level range of [0..4095] into k_1 equally sized bins, while *clipped binning* allocates one large bin for low intensity gray-levels [0..855] (air), one for higher intensities [1368..4095] (bone), and k_2 equally sized bins for the soft tissues between [856..1368]. *Nonlinear binning* divides the gray-level range of [0..4095] into k_3 bins of different sizes. These bins are further used to calculate the co-occurrence statistical model and its ten Haralick descriptors for texture quantification of gray-level images. The results of the texture quantification using each one of the three strategies and for different values of k_1 , k_2 and k_3 are evaluated with respect to their discrimination power using a decision tree classification algorithm and four classification performance metrics (sensitivity, specificity, precision and accuracy). Our preliminary results obtained on 1368 segmented DICOM images show that the optimal number of gray-levels is equal to 128 for linear binning, 512 for clipped binning, and 256 for non-linear binning. Furthermore, when comparing the results of the three approaches, the *nonlinear binning approach shows significant improvement for heart and spleen*.

Keywords: Binning, texture, co-occurrence matrix, computed tomography

1. INTRODUCTION

Binning is a term used to describe the reduction of gray level intensities in an image, in this case by combining several intensity levels into a single intensity level, or bin. *Linear binning* describes the process of equally dividing the original gray level range of the image into k bins. *Clipped binning* describes the process of determining a range of interest, where the gray levels are equally divided into k bins, while the intensities above the desired range are placed into the highest bin, and the intensities below the desired range are placed into the lowest bin. *Nonlinear binning* describes the process of using histogram data to assign the original gray level range into possibly unequal bins.

In this paper we evaluate three types of gray-level binning strategies, *linear binning*, *clipped binning*, and *nonlinear binning*, applied for efficient and accurate texture characterization of normal tissues in Computed Tomography (CT) of the chest and abdomen. The gray-level binning strategies serve as an image preprocessing step for the computation of co-occurrence matrices and Haralick texture features. The objective of this paper is to identify the binning method that reduces the gray-level resolution and the data redundancy present in the image while optimizing the classification

[♦] This material is based upon work supported by the National Science Foundation under Grant No. 0453456.

accuracy of five healthy tissues (backbone, liver, heart, renal, and splenic parenchyma) based on ten Haralick texture descriptors.

2. BACKGROUND

Previous work on the topic of binning has been done to determine an optimal k value using entropy and information gain⁸. Entropy is a measure of the information content of an image, and can be used to determine how many gray levels are necessary to represent all of the information¹¹. Other work in the medical field evaluates a type of histogram binning compared to linear binning with respect to the speed and accuracy of pixel co-registration in magnetic resonance imaging (MRI)¹. A method of classifying different lung disease as well as normal lung tissue utilized 16 gray levels but gave no explanation for choosing this particular k value².

3. METHODOLOGY

3.1 The data set

Our data set consisted of 1368 segmented DICOM images. There were 141 chest and abdomen high resolution CT images from two healthy patients. The five organs (backbone, liver, heart, renal, and splenic parenchyma) were then segmented from each of these CT images, resulting in around 340 segmented images for the group of organs. In order to extract more information from each organ and ultimately build a better classifier, each of the images was divided into quadrants resulting in the 1376 images. Eight images were removed because too much background data was present in these images and the texture data were not useful. The tissues represented in these images are all normal tissues, because if normal tissues can be classified correctly, then abnormal tissues would be easily identified.

3.2 Definition of clipped, linear, and nonlinear binning

Linear binning divides the gray-level range of [0..4095] into k_1 equally sized bins, *clipped binning* allocates one large bin for low intensity gray-levels [0..855] (air), one for higher intensities [1368..4095] (bone), and k_2 equally sized bins for the soft tissues between [856..1368], and *nonlinear binning* distributes the gray levels into unequal bins based on histogram data. All types of binning are evaluated for five different k values: 32, 64, 128, 256, and 512; for nonlinear binning, this k value represents the number of clusters to be used in the k -means clustering⁹ approach, and does not represent the final number of gray levels in the image⁸.

3.3 Definition of the Co-occurrence Matrix

The co-occurrence matrix is often used in texture analysis. By specifying a displacement vector, \mathbf{d} , and counting all pairs of pixels separated by \mathbf{d} with gray level intensities i and j , this co-occurrence matrix is created⁶. An image with 512 gray levels will have a 512 x 512 co-occurrence matrix, while an image with 32 gray levels will have a co-occurrence matrix that is only 32 x 32. This large size difference can have a huge impact on efficiency and robustness when using these matrices to calculate texture, thus the need to reduce the gray levels in the image, through binning, before analyzing texture.

In the case of this study, four different directions and five different distances were used to create the displacement vector \mathbf{d} . The distances range from 1 to 5 pixels, while the directions are 0° , 45° , 90° , and 135° . In total, there are 20 possible displacement vectors, resulting in 20 co-occurrence matrices to describe each image. Different displacement vectors result in different co-occurrence matrices; this can be seen in Figure 1, where a co-occurrence matrix is calculated for a 5 x 5 image for distance 1, direction 90° , and distance 1, direction 135° .




Figure 1: (a) 5 x 5 image with 4 gray levels. (b) Co-occurrence matrix for $\mathbf{d} = (1,0)$. (c) Co-occurrence matrix for $\mathbf{d} = (1,1)$.

3.4 Co-occurrence Features

There are many different numerical calculations that can be computed from the co-occurrence matrix. The ten descriptors used in this work, as described by Haralick⁵, are as follows.

3.4.1 Entropy

$$-\sum_i^M \sum_j^N P[i, j] \log P[i, j]$$

3.4.2 Energy

$$\sum_i^M \sum_j^N P^2[i, j]$$

3.4.3 Contrast

$$\sum_i^M \sum_j^N (i - j)^2 P[i, j]$$

3.4.4 Homogeneity

$$\sum_i^M \sum_j^N \frac{P[i, j]}{1 + |i - j|}$$

3.4.5 SumMean

$$\frac{1}{2} \sum_i^M \sum_j^N (iP[i, j] + jP[i, j])$$

3.4.6 Variance

$$\frac{1}{2} \sum_i^M \sum_j^N ((i - \mu)^2 P[i, j] + (j - \mu)^2 P[i, j])$$

3.4.7 Correlation

$$\sum_i^M \sum_j^N \frac{(i - \mu)(j - \mu)P[i, j]}{\sigma^2}$$

3.4.8 Maximum Probability

$$\overline{Max}_{i,j} P[i, j]$$

3.4.9 Inverse Difference Moment

$$\sum_i^M \sum_j^N \frac{P[i, j]}{|i - j|^k} \quad i \neq j$$

3.4.10 Cluster Tendency

$$\sum_i^M \sum_j^N (i + j - 2\mu)^k P[i, j]$$

These particular texture features were found to work well in the previous work done on tissue classification³. After the texture descriptors are calculated, each image in the data set is represented by a feature vector [D₁, D₂... D₁₀, class label] consisting of 10 texture descriptors and a class label annotating the corresponding region; in the current implementation, the label takes one of the following values: ‘backbone’, ‘heart’, ‘spleen’, ‘kidney’, or ‘liver’. These texture descriptors are calculated for each co-occurrence matrix, so for each *k* value and type of binning there are 20 different sets of feature vectors for the data set of images. In order to summarize the data, each texture descriptor is averaged over all 20 values for each *k*, and then normalized between 0 and 1 using a min-max criteria⁴ to bring all of the descriptor data into the same scale. A diagram of this procedure can be seen in Figure 2.




Figure 2: Averaging texture descriptor data.

3.5 Decision Tree Classification

To classify the images in the data set, we divided the entire data set into a training set containing a random sampling of 66% of the original data and an ‘unseen’ data set containing the other 34% of the original data. The training set is used to build the tree, while the second set, the ‘unseen’ set, is used to estimate the accuracy of the tree when used for classification of unseen images.

To avoid over-fitting the tree to the data set and producing an overoptimistic estimation of the classifier accuracy, a *k*-fold cross-validation technique was used to train each tree. The training data is randomly divided into *k* mutually exclusive ‘folds’ (S₁, S₂ . . . S_k), each with similar class proportions and of approximately the same size. Training and testing of the tree are performed *k* times, and for each iteration *i*, the subset S_{*i*} is reserved as the test set, while the remaining sets are used to train the tree. The accuracy is estimated from the overall number of correct classifications from the *k* iterations divided by the total number of samples⁴. In our results, we use 10-fold cross-validation for estimating the classifier accuracy.

3.6 C&RT decision trees

There are many types of classifiers that could be used to identify organ tissues from a data set. Decision trees have a relatively fast learning speed, while still producing good classification accuracy, and generate decision rules that are easily understood and applied.

The trees were implemented using a Classification and Regression Tree (C&RT) approach, which chooses a predictor, in this case a texture descriptor, to split the current node into two children. For each split, each descriptor is evaluated to find the best cut point with which to describe the children, the goal being to produce subsets, or leaves, that are as pure as possible with respect to the class label. There are several conditions that will stop tree growth³:

- the maximum tree depth, d , has been reached,
- there is no significant predictive descriptor left to split the node,
- the number of cases in the terminal node is less than the minimum number, n_p , of cases for parent nodes,
- if the terminal node were to split, the number of cases in one or more child nodes would be less than the minimum number, n_c , of cases for child nodes,
- minimum change in impurity, imp , is reached.

Depending on the parameter values (d , n_p , n_c , imp) a different tree will be produced. To find the best tree, the depth parameter was fixed to 20, the impurity to 0.0001, and the number of child nodes to 1. The number of parents was increased, starting from 2, until the sensitivity and specificity of the testing set reached a maximum (these performance metrics are discussed below). After finding this optimal parent, the number of children was increased from 1 to $n_p/2$. The complexity of the tree then decreases, and the performance increases on the ‘unseen’ data.

3.7 Performance Metrics

In order to select the ‘best’ parent and ‘best’ child, and so the ‘best’ decision tree for our data, we use four metrics to measure the performance of the classifier; the tree which will maximize the values of all four metrics will be defined as the ‘best’ tree.

We define these metrics in terms of positives and negatives. As an example, assume that the positives are the heart tissues and the negatives are the non-heart tissues. A true positive is a tissue region classified as heart when the original class label (the label given by a human expert) is heart, true negative is a tissue region correctly classified as non-heart, false positive is a tissue region classified as heart when it is actually a non-heart, and a false negative is a tissue region classified as non-heart when it is actually a heart. The four metrics¹⁰ can be seen below, in Table 1.

Metric	Description
Sensitivity	true positives/(true positives + false negatives)
Specificity	true negatives/(true negatives + false positives)
Precision	true positives/(true positives + false positives)
Accuracy	(true positives + true negatives)/total samples

Table 1: Description of performance metrics.

Based on the performance metrics of overall sensitivity and specificity, the best tree for each k value was determined to be the tree with the highest of these values. After finding the optimal tree for each k value, results within each binning type are compared to find an optimal tree and optimal k value for linear and clipped binning. The optimal linear, clipped, and non-linear results are then compared to find the optimal binning method with respect to classification.

4. RESULTS

We evaluated the three binning strategies on 1368 tissue images segmented from 141 DICOM images of two normal high resolution CT studies. The tables of optimal performance metrics for each k and each binning method can be seen below, in Tables 2-16, where Tables 2-6 represent linear binning, Tables 7-11 represent clipped binning, and Tables 12-16 represent nonlinear binning.

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	78.873%	97.949%	76.056%	72.464%	36.957%	81.416%
Specificity	97.375%	96.498%	95.538%	91.645%	96.059%	95.700%
Precision	84.848%	95.500%	76.056%	60.976%	51.515%	81.026%
Accuracy	94.469%	97.124%	92.478%	88.717%	90.044%	93.973%

Table 2: Performance metrics for linear binning, $k = 32$.

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	88.732%	97.436%	85.915%	69.565%	60.870%	86.283%
Specificity	96.325%	97.665%	96.850%	96.606%	95.813%	96.977%
Precision	81.818%	96.939%	83.562%	78.689%	62.222%	86.143%
Accuracy	95.133%	97.566%	95.133%	92.478%	92.257%	95.485%

Table 3: Performance metrics for linear binning, $k = 64$.

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	88.000%	97.449%	78.261%	73.973%	65.306%	85.931%
Specificity	97.158%	98.872%	96.947%	95.630%	94.673%	97.349%
Precision	85.714%	98.454%	81.818%	76.056%	59.259%	86.205%
Accuracy	95.671%	98.268%	94.156%	92.208%	91.558%	95.563%

Table 4: Performance metrics for linear binning, $k = 128$.

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	88.571%	98.985%	80.282%	61.250%	59.259%	83.686%
Specificity	94.776%	99.636%	97.007%	95.408%	94.019%	97.161%
Precision	74.699%	99.490%	82.609%	73.134%	56.140%	83.847%
Accuracy	93.856%	99.364%	94.492%	89.619%	90.042%	95.096%

Table 5: Performance metrics for linear binning, $k = 256$.

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	89.744%	98.930%	84.848%	66.216%	57.895%	85.065%
Specificity	96.615%	98.182%	95.455%	95.103%	96.543%	96.832%
Precision	84.337%	97.368%	75.676%	72.059%	70.213%	84.665%
Accuracy	95.455%	98.485%	93.939%	90.476%	91.775%	95.213%

Table 6: Performance metrics for linear binning, $k = 512$.

	Kidney	Backbone	Heart	Liver	Spleen	Overall
--	--------	----------	-------	-------	--------	---------

Sensitivity	88.732%	95.897%	80.282%	82.609%	54.348%	86.062%
Specificity	98.688%	96.498%	95.538%	94.256%	97.537%	96.455%
Precision	92.647%	95.408%	77.027%	72.152%	71.429%	86.097%
Accuracy	97.124%	96.239%	93.142%	92.478%	93.142%	95.002%

Table 7: Performance metrics for clipped binning, $k = 32$.

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	79.747%	96.500%	80.769%	78.947%	54.386%	83.673%
Specificity	96.594%	98.621%	96.359%	93.720%	95.150%	96.770%
Precision	81.818%	97.970%	80.769%	69.767%	59.615%	83.792%
Accuracy	93.878%	97.755%	93.878%	91.429%	90.408%	94.677%

Table 8: Performance metrics for clipped binning, $k = 64$.

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	91.549%	97.949%	85.915%	71.014%	58.333%	86.052%
Specificity	96.456%	99.262%	96.456%	95.466%	95.813%	97.401%
Precision	82.278%	98.964%	81.333%	73.134%	67.308%	85.835%
Accuracy	95.708%	98.712%	94.850%	91.845%	90.987%	95.655%

Table 9: Performance metrics for clipped binning, $k = 128$.

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	88.732%	97.740%	80.000%	71.250%	50.000%	82.432%
Specificity	94.906%	97.753%	95.939%	94.780%	95.238%	96.184%
Precision	76.829%	96.648%	71.429%	75.000%	64.706%	81.990%
Accuracy	93.919%	97.748%	94.144%	90.541%	88.514%	94.058%

Table 10: Performance metrics for clipped binning, $k = 256$.

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	87.324%	98.974%	78.873%	79.710%	65.217%	87.611%
Specificity	96.588%	97.665%	98.950%	96.606%	95.074%	97.272%
Precision	82.667%	96.985%	93.333%	80.882%	60.000%	87.940%
Accuracy	95.133%	98.230%	95.796%	94.027%	92.035%	96.089%

Table 11: Performance metrics for clipped binning, $k = 512$.

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	80.000%	91.795%	78.261%	72.857%	52.174%	80.889%
Specificity	96.053%	92.157%	95.801%	97.105%	94.059%	94.286%
Precision	78.873%	89.950%	77.143%	82.258%	50.000%	80.983%
Accuracy	93.556%	92.000%	93.111%	93.333%	89.778%	92.393%

Table 12: Performance metrics for nonlinear binning, $k = 32$ (41 GL).

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	76.364%	91.667%	72.222%	89.412%	56.364%	82.378%
Specificity	96.875%	93.258%	97.243%	93.523%	96.154%	94.676%
Precision	76.364%	91.220%	82.540%	75.248%	65.957%	82.326%

Accuracy	94.480%	92.569%	93.418%	92.781%	91.507%	92.836%
----------	---------	---------	---------	---------	---------	---------

Table 13: Performance metrics for nonlinear binning, $k = 64$ (65 GL).

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	91.250%	87.179%	82.090%	73.494%	60.377%	81.799%
Specificity	95.226%	96.113%	95.377%	96.456%	94.353%	95.726%
Precision	79.348%	93.923%	74.324%	81.333%	57.143%	82.472%
Accuracy	94.561%	92.469%	93.515%	92.469%	90.586%	92.757%

Table 14: Performance metrics for nonlinear binning, $k = 128$ (70 GL).

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	87.143%	94.872%	85.507%	78.571%	73.913%	87.556%
Specificity	97.368%	96.863%	97.900%	97.895%	94.554%	97.025%
Precision	85.915%	95.855%	88.060%	87.302%	60.714%	88.191%
Accuracy	95.778%	96.000%	96.000%	94.889%	92.444%	95.429%

Table 15: Performance metrics for nonlinear binning, $k = 256$ (68 GL).

	Kidney	Backbone	Heart	Liver	Spleen	Overall
Sensitivity	92.754%	96.907%	79.710%	88.406%	47.826%	87.248%
Specificity	96.561%	98.024%	98.677%	94.974%	96.259%	97.246%
Precision	83.117%	97.409%	91.667%	76.250%	59.459%	87.145%
Accuracy	95.973%	97.539%	95.749%	93.960%	91.275%	95.824%

Table 16: Performance metrics for nonlinear binning, $k = 512$ (50 GL).

For each type of binning, one particular k value performed better than all the others. In the case of linear binning, $k = 128$ performed the best, with tree parameters of 12 and 2 for parent and child, respectively. For clipped binning, $k = 512$ had the highest sensitivity and specificity, with a parent node of 22 and a child node of 6. For nonlinear binning, $k = 256$ performed the best, with a parent node of 16 and a child node of 1.

At the organ level, different k values represented optimal performance. For each type of binning, the optimal k value for each organ, as well as the parent and child node parameters, can be see in Tables 17 -19.

Organ	Optimal k	Parent node	Child node	Sensitivity	Specificity
Kidney	512	9	5	91.026%	95.833%
Backbone	256	3	1	99.492%	98.909%
Heart	64	24	1	87.324%	96.588%
Liver	128	35	1	82.192%	93.316%
Spleen	128	35	1	67.347%	95.400%

Table 17: Optimal k values for each organ, linear binning.

Organ	Optimal k	Parent node	Child node	Sensitivity	Specificity
Kidney	128	6	1	92.958%	96.203%
Backbone	128	24	1	98.974%	99.262%
Heart	128	18	1	85.915%	96.709%
Liver	256	55	1	80.000%	91.209%
Spleen	512	22	6	65.217%	95.074%

Table 18: Optimal k values for each organ, clipped binning.

Organ	Optimal k	Parent node	Child node	Sensitivity	Specificity
Kidney	512	16	1	91.429%	97.105%
Backbone	512	10	2	96.392%	98.039%
Heart	256	23	1	86.957%	97.638%
Liver	512	12	1	90.000%	94.474%
Spleen	256	16	3	76.087%	94.059%

Table 19: Optimal k values for each organ, nonlinear binning.

A useful tool for visualizing the most effective tree for a particular k value or the most effective tree overall is the Receiver Operating Characteristic (ROC) curve, which plots sensitivity on the y-axis and 1-specificity on the x-axis. A point in the upper left corner of the graph represents an optimal tree. The ROC curve plotting the performance of the different trees tested for linear binning, $k = 128$ can be seen in Figure 3. The red point represents the optimal tree with parent node 12 and child node 2.

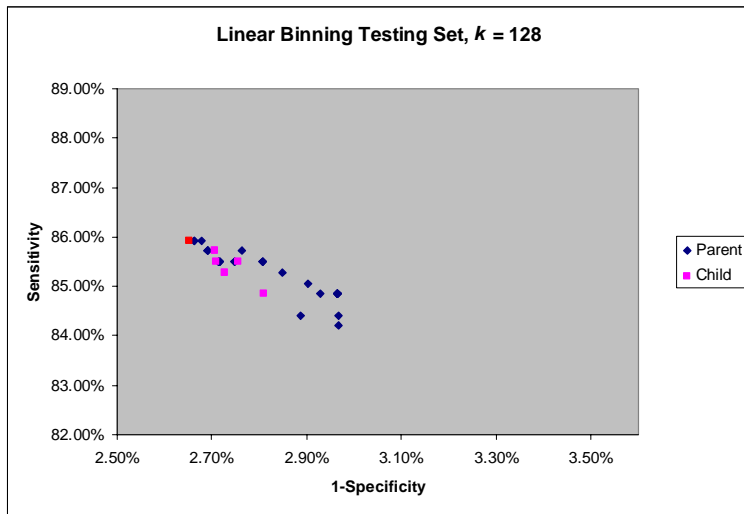


Figure 3: ROC curve for optimal trees for each k value, linear binning.

The ROC curve plotting the performance of the different trees tested for clipped binning, $k = 512$ can be seen in Figure 4. The red point represents the tree for parent node 22, child node 6.

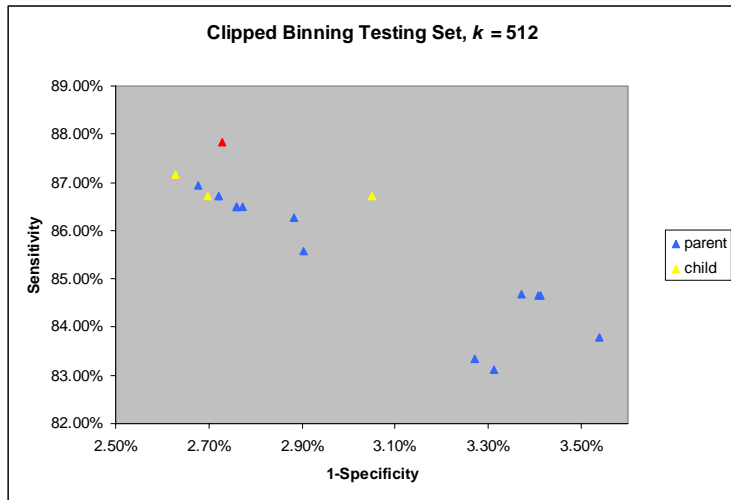


Figure 4: ROC curve for clipped binning, $k = 64$.

The ROC curve plotting the performance of the different trees tested for nonlinear binning, $k = 256$ can be seen in Figure 5. The red dot represents the optimal tree with parent node 16, child node 1.

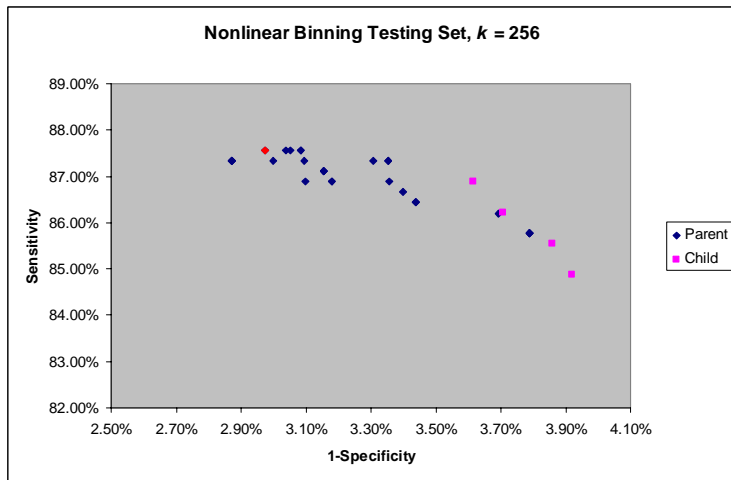


Figure 5: ROC curve for nonlinear binning, $k = 256$.

The tables each utilize the same x and y scales, so it is easy to see that both clipped and nonlinear binning outperform linear binning with respect to classification. From the graphs, clipped appears to outperform nonlinear binning with a slightly higher overall sensitivity and specificity.

5. CONCLUSIONS

For linear binning, $k = 128$ produced the best overall sensitivity (85.93%) and specificity (97.35%); for clipped binning, $k = 512$ produced the best overall sensitivity (87.61%) and specificity (97.27%). For nonlinear binning, $k = 256$ produced the best overall sensitivity (87.56%) and specificity (97.03%). The difference between the overall sensitivities and specificities for linear and clipped, or linear and nonlinear, are large enough to see that linear binning is not the optimal method to use in order to improve classification results. When comparing the sensitivity and specificity of nonlinear and clipped, clipped has the higher overall result, but in the case of nonlinear, the sensitivity for the spleen was

7% higher than that of clipped. The heart sensitivity is also improved by nonlinear binning, increasing to 85.51% from 78.87%. The other organs, kidney, backbone, and liver, have comparable sensitivity values for nonlinear and clipped. A larger number of gray levels, in both nonlinear and clipped, was used to represent the soft tissue range; for clipped, 510 gray levels represent the soft tissues while only 16 represent the soft tissues in linear. For nonlinear, 68 gray levels are used to display the entire image, divided based on gray level occurrences. Though clipped binning may have more gray levels to represent the soft tissue range than nonlinear binning, nonlinear is based on the actual content of the image and so may be better representing the data. Nonlinear binning turned out to be the best binning method, which was what was initially hypothesized. This is a logical result, because the bins allocated to the soft tissue range represent actual content present in the image.

6. BIOGRAPHIES

Bahare Naimipour is a senior at the University of Illinois at Chicago. She majors in Electrical Engineering, and minors in Mathematics. Her fields of interest include Signal and Image Processing, and Applied Mathematics. She is currently the Vice President of The Professional Engineers Society Counsel, Society of Women Engineers officer, an active member of the Honors College, and has recently joined ETA KAPPA NU along with IEEE. A native of Chicago, Bahare knows four languages, of which two she can speak and write fluently. Her hobbies include tutoring math and landscape photography.

Stefanie Handrick is a junior at Arizona State University in Tempe, Arizona. She is a Computer Systems Engineering major, with interests including reading, web design, and assembly language programming. She is a member of Kappa Kappa Psi, the national band fraternity, and has been awarded several scholarships, including one from Intel, SWE, and Tektronix. She is a Phoenix native, whose hobbies include learning to speak Spanish.

7. ACKNOWLEDGEMENTS

This work was made possible by NSF grant number IIS-0453456.

8. REFERENCES

- [1] J. Camp, R. Robb, "A novel binning method for improved accuracy and speed of volume image coregistration using normalized mutual information", **Proc. SPIE – The International Society for Optical Engineering**, 3661:24-31, 1999.
- [2] F. Chabat, G. Yang, D. M. Hansell, "Obstructive Lung Diseases: Texture Classification for Differentiation at CT", **Radiology** **2003**, 228:871-877.
- [3] D. Channin, D. S. Raicu, J. D. Furst, D. H. Xu, L. Lilly, & C. Limpsangsri, "Classification of Tissues in Computed Tomography using Decision Trees", Poster and Demo, **The 90th Scientific Assembly and Annual Meeting of Radiology Society of North America (RSNA04)**, Chicago, IL, USA, November 28 - December 3rd, 2004.
- [4] J. Han & M. Kamber, *Data Mining: Concepts and Techniques*. Morgan. Kaufmann Publishers, 2001.
- [5] R. M. Haralick, "Statistical and structural approaches to texture," **Proc. IEEE**, vol. 67, pp. 786--804, 1979.
- [6] R. Jain, R. Kasturi, & B. G. Schunck, *Machine Vision*. New York: McGraw-Hill, 1995.
- [7] A. Kurani, D. H. Xu, J. D. 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, in August 16-18, 2004.
- [8] R. Lerman, D. Raicu, J. Furst. "Contrast enhancement of soft tissues in Computed Tomography images," **SPIE Medical Imaging, February 2006**.

[9] R. J. Roiger & M. W. Geatz, *Data Mining: A Tutorial-Based Primer*. Addison Wesley, 2003.

[10] L. Semler, L. Dettori, J. Furst, "*Wavelet-Based Texture Classification of Tissues in Computed Tomography*", **The 18th IEEE International Symposium on Computer-Based Medical Systems(CBMS'05)**, Dublin, Ireland, June 23-24, 2005.

[11] M. Sonka, V. Hlavac, & R. Boyle, *Image Processing, Analysis, and Machine Vision*. PWS Publishing, 1998.