

# **A Similarity Learning Approach to Content Based Image Retrieval: Application to Digital Mammography**

Issam El-Naqa, *Student Member, IEEE*, Yongyi Yang\*, *Senior Member, IEEE*,  
Nikolas P. Galatsanos, *Senior Member, IEEE*, Robert M. Nishikawa, and  
Miles N. Wernick, *Senior Member, IEEE*,

This work was supported by NIH/NCI grant CA89668.

*Asterisk denotes corresponding author.*

I. El-Naqa is with the Medical School of Washington University, 660 S. Euclid, St. Louis, MO 63110.

Y. Yang, M. N. Wernick, and N. P. Galatsanos are with the Department of Electrical and Computer Engineering, Illinois Institute of Technology, 3301 South Dearborn Street, Chicago, IL 60616.

R. M. Nishikawa is with the Department of Radiology, The University of Chicago, 5841 South Maryland Avenue, Chicago, IL 60637.

## ABSTRACT

In this paper we describe an approach to content-based retrieval of medical images from a database, and provide a preliminary demonstration of our approach as applied to retrieval of digital mammograms. Content-based image retrieval (CBIR) refers to the retrieval of images from a database using information derived from the images themselves, rather than solely from accompanying text indices. In the medical-imaging context, the ultimate aim of CBIR is to provide radiologists with a diagnostic aid in the form of a display of relevant past cases, along with proven pathology and other suitable information. CBIR may also be useful as a training tool for medical students and residents.

The goal of information retrieval is to recall from a database information that is relevant to the user's query. The most challenging aspect of CBIR is the definition of relevance (similarity), which is used to guide the retrieval machine. In this paper, we pursue a new approach, in which similarity is learned from training examples provided by human observers. Specifically, we explore the use of neural networks and support vector machines to predict the user's notion of similarity. Within this framework we propose using a hierarchical learning approach, which consists of a cascade of a binary classifier and a regression module to optimize retrieval effectiveness and efficiency. We also explore how to incorporate online human interaction to achieve relevance feedback in this learning framework. Our experiments are based on a database consisting of 76 mammograms, all of which contain clustered microcalcifications (MCs). Our goal is to retrieve mammogram images containing similar MC clusters to that in a query. The performance of the retrieval system is evaluated using precision-recall curves computed using a cross-validation procedure. Our experimental results demonstrate that: (1) the learning framework can accurately predict the perceptual similarity reported by human observers, thereby serving as a basis for CBIR; (2) the learning-based framework can significantly outperform a simple distance-based similarity metric; (3) the use of the hierarchical two-stage network can improve retrieval performance; and (4) relevance feedback can be effectively incorporated into this learning framework to achieve improvement in retrieval precision based on online interaction with users.

*Index Terms*—Content-based image retrieval, kernel methods, mammogram, digital radiography, computer-aided diagnosis, relevance feedback.

## I. INTRODUCTION

Content-based image retrieval (CBIR) refers to the recall of relevant images to a query from a database, using information derived from the images themselves, rather than relying on accompanying text indices or other annotation. CBIR has received increasing attention as a result of the availability of large image databases in medicine, science, commerce, and the military [1,2]. CBIR has been proposed to overcome the difficulties encountered in textual annotation for large image databases. Like a text-based search engine, a CBIR system aims to retrieve information that is relevant (or similar) to the user's query. In document retrieval, the query is usually a word or phrase; in CBIR, it is an image. The key to successful CBIR lies in the development of appropriate similarity metrics for ranking the relevance of images in a database to the query image. In CBIR quantitative image features, computed automatically, are used to characterize image content. The image features may be extracted at either a low level (such as local edges [3]) or at a high level (such as color histogram [4]), or both. The query image is then compared to the images in the database on the basis of the measured features. Those images in the database having the highest similarity to the query image are retrieved and displayed for the user.

The general application of image retrieval to broad image databases has experienced limited success, principally due to the difficulty of quantifying image similarity for unconstrained image classes (e.g., all images on the Internet). We expect that medical imaging will be an ideal application of CBIR, because of the more-limited definition of image classes (e.g., digital mammograms), and because the meaning and interpretation of medical images is better understood and characterized. In spite of this, the application of CBIR in medical imaging thus far has been somewhat limited [5]. In [6], a rule-based expert system was developed to display chest radiographs from a library of images as illustrative examples for helping radiologists' diagnosis. In [7], a retrieval method based on texture and shape analyses was applied for

search and retrieval of a database containing pulmonary CT images. In [8], an algorithm was described for retrieval of 3D MR images based on anatomical structure matching. In [9], a similarity metric based on Bayes decision theory was developed for retrieval of neuroradiological CT images. In [10,11], a technique was developed that reduces high-dimensional data to a two-dimensional feature space in which images that are close to each other are selected for purposes of visualizing relationships in the data. In [12], a retrieval method was developed using correlation coefficients in a database of pulmonary nodules represented by the joint histogram of the pattern CT density and 3-D curvature shape index.

### *A. A Learning Approach to Quantify Image Similarity*

Unlike the existing approaches to CBIR, which are typically based on some simple distance measures for image similarity, we propose an approach in which machine-learning algorithms (neural networks and support vector machines) are trained to predict the measures of image similarity reported by human observers. We treat the learning of the similarity function as a nonlinear regression of the similarity coefficient on the features of the images. The method is developed using a set of 76 mammograms, all containing clustered microcalcifications (MCs). Our goal is to retrieve mammograms containing similar MC clusters to that in a query mammogram. The proposed retrieval framework is evaluated statistically using a cross-validation procedure.

The feasibility of a learning-based approach for modeling perceptual similarity was first demonstrated in our previous work in [13] using simulated image data. In this work, we expand this approach in two major aspects. First, we develop a hierarchical two-stage learning network for improved performance. Second, we explore how to utilize user interaction, known as relevance feedback, in the learning framework so as to achieve online adaptation to the user.

### *B. Application to Mammography*

Mammography has been by far the most effective means for early detection of breast cancer, a leading cause of death in women in many developed countries. The sensitivity of mammography is approximately 90% [14]. In spite of the technological advances in recent years, mammogram reading still

remains a difficult clinical task. Some breast cancers may produce changes in mammograms that are subtle and difficult to recognize. It has been reported that 10-30% of lesions are misinterpreted during routine screening of mammograms [15]. Furthermore, it is very difficult to distinguish benign lesions from malignant ones in mammograms. As a result, between 2 and 10 women are biopsied for every cancer detected, causing needless fear and pain to women who are biopsied [16,17]. This low specificity results in a relatively large inter-observer variability that can lead to failure to biopsy malignant lesions and potentially avoidable biopsy of benign lesions [18].

We conjecture that by presenting images with known pathology that are “visually similar” to the image being evaluated, the use of a mammogram retrieval system may provide a more intuitive aid to radiologists, potentially leading to improvement in their diagnostic accuracy. Furthermore, it is expected that the proposed technique would be a useful aid in the training of students and residents, since it would allow them to view images of lesions that appear similar, but may have differing pathology.

An alternative approach to computer-aided diagnosis (CAD), in which the likelihood of malignancy is computed (e.g., [19]), has been studied to a large extent in the literature. The proposed retrieval system is in principle very different, and may helpfully complement existing diagnostic aids. Our retrieval system follows a “critiquing” approach [20]: instead of proposing a diagnosis, it aims to assist the radiologist by providing relevant supporting evidence from prior known cases. If we view the human observer as a classifier, then the aim of the CBIR system is to provide the observer with training-set examples that are close to his decision boundary, along with the correct class labels (proven pathology) for these examples. The hypothesis that we ultimately hope to demonstrate is that this approach will improve the classification (diagnostic) performance of the observer.

In developing a CBIR system for digital mammography, we argue that the similarity metric must conform closely to the user’s notions of similarity, and that simple, mathematical distance metrics may not be adequate for describing perceived similarity. Therefore, we aim to show that a learned concept of similarity can outperform simple distance metrics in modeling the user’s similarity concept.

The remainder of the paper is organized as follows. First, an overview of the proposed learning-based retrieval framework is provided in Section II. In Section III the hierarchical learning network is described, and relevance-feedback techniques are developed in Section IV. An evaluation study, including data-set acquisition, training, and testing procedures is described in Section V. Experimental results are presented in Section VI. Finally, conclusions are drawn in Section VII.

## II. OVERVIEW OF THE PROPOSED IMAGE-RETRIEVAL FRAMEWORK

We assume that the user's notion of similarity between a pair of images is a function of the relevant features in the images. We then use machine learning to model this notion of similarity for the purpose of CBIR. Our goal is to find, among the many images in the database, those that are most visually similar to the query as judged by the user.

The proposed framework is illustrated with a functional diagram in Fig. 1. For a given query image, we first characterize its content by an  $M$ -dimensional vector  $\mathbf{u}$ , quantifying the key relevant features of the image. This feature vector is then compared to the corresponding feature vector  $\mathbf{v}$  of a database entry by way of a learning machine, denoted by a nonlinear mapping  $f(\mathbf{u}, \mathbf{v})$ , to produce a similarity coefficient ( $SC$ ). The images with the highest  $SC$ s (say, those above a prescribed threshold value  $T$ ) are then retrieved from the database.

Clearly, the key to this framework<sup>1</sup> lies in the learning machine  $f(\mathbf{u}, \mathbf{v})$ . Ideally, the learning machine should have the following properties: 1)  $f(\mathbf{u}, \mathbf{v})$  must closely conform to the user's notion of similarity; 2) the learning machine should involve reasonable computational complexity so that it can be applied to a large-scale database; and 3)  $f(\mathbf{u}, \mathbf{v})$  should provide the user with the ability to refine the search in a process called relevance feedback (indicated by the dashed path in Fig. 1).

We adopt a supervised learning approach for  $f(\mathbf{u}, \mathbf{v})$ . For this purpose we first collect a set of sample image pairs, each having a labeled  $SC$  (e.g., obtained from observer studies). We then train a learning

<sup>1</sup> Of course, an equally important issue, if not more important, is the selection of feature vector  $\mathbf{u}$  so that its components are relevant to perceptual similarity. In this study, we will make use of some features already existing in the literature (detail in Section V.C).

machine  $f(\mathbf{u}, \mathbf{v})$  with these samples. Specifically, letting  $SC(\mathbf{u}, \mathbf{v})$  denote the similarity coefficient between an image pair  $\mathbf{u}$  and  $\mathbf{v}$ , we model  $SC(\mathbf{u}, \mathbf{v})$  as:

$$SC(\mathbf{u}, \mathbf{v}) = f(\mathbf{u}, \mathbf{v}) + \xi, \quad (1)$$

where  $\xi$  is the modeling error of the learning machine. Our aim is to determine a functional  $f(\cdot, \cdot)$  that will generalize well to images outside the training set.

Since our aim is always the comparison of pairs of images, we will view the similarity metric as a functional of a single argument  $\mathbf{x} \equiv \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix}$ , which is a concatenation of the feature vectors  $\mathbf{u}$  and  $\mathbf{v}$  of two images to be compared; thus, we redefine the similarity functional  $f(\mathbf{u}, \mathbf{v})$  as  $f(\mathbf{x})$ .

### III. HIERARCHICAL LEARNING NETWORK

We propose to use a two-stage hierarchical learning network to model the perceptual similarity for retrieval. This network consists of a cascade of a binary classifier stage and a regression stage for predicting the  $SC$ s between a query image and the images in the database, as illustrated in Fig. 2. In the first stage, images that are very different from the query image are eliminated from further consideration by a binary classifier. Images surviving this stage are then compared to the query in the second stage to obtain a numerical  $SC$  for retrieval.

The learning network in Fig. 2 is hierarchical in the following sense: during the training phase, the first-stage classifier functions as a coarse, binary learner, the purpose of which is for triage; i.e., the first stage identifies simply whether a database entry is sufficiently similar to the query for further consideration. The second stage functions as a more-refined learner, the purpose of which is to measure quantitatively the similarity between a surviving database entry and the query.

The reasons for this approach are as follows. First, the triage classifier avoids the computational cost of carefully measuring the  $SC$ s of those database entries that are not at all similar to the query, and thus will certainly not be determined to be relevant. Second, by training the second stage only using

reasonably similar pairs, the learning machine can be better fine-tuned to predict *SCs* for those image pairs that are of genuine interest.

Of course, the use of a triage stage can also have adverse effect, i.e., it may eliminate some truly similar images from further consideration. To ameliorate this effect, in the following we will modify the cost function of in the SVM classifier such that it will impose a greater penalty on missed similar images than on misclassified non-similar images. As demonstrated by our experimental results (Section VI), this approach can lead to significant improvement in retrieval performance. Below we discuss the details of the two-stage network.

#### A. First-Stage Classifier

Consider a query image and a database entry with feature vectors  $\mathbf{u}$  and  $\mathbf{v}$ , respectively. The task of the first-stage classifier is to determine whether the two images are sufficiently similar for further consideration. This is treated as a two-class pattern classification problem, i.e., the mammogram image pair is either reasonably similar (designated as “class 1”) or not similar (designated as “class 2”). For reasons of computational speed, we employ a linear classifier for this task, i.e., we use a decision function of the form

$$h(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + b \quad (2)$$

such that  $h(\mathbf{x}) \geq 0$  if the image pair  $\mathbf{x}$  is sufficiently similar, and  $h(\mathbf{x}) < 0$  otherwise. In other words, image pairs from the two different classes are separated by the hyperplane  $h(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + b = 0$ .

The decision function  $h(\mathbf{x})$  is to be determined from training samples. Let  $\{(\mathbf{x}_i, d_i), i = 1, 2, \dots, l\}$  denote a given set of  $l$  training examples, where each sample pair  $\mathbf{x}_i$  has a known class label  $d_i$  (i.e.,  $d_i = +1$  for class 1, and  $d_i = -1$  for class 2). The problem then is how to determine  $\mathbf{w}$  and  $b$  in  $h(\mathbf{x})$  so that it that can correctly classify an input pattern (not necessarily from the training set).

We consider two types of pattern classifiers for this task: (1) a Fisher discriminant, and (2) a support vector machine (SVM). To distinguish between the two classifiers, below we use  $\underline{\mathbf{w}}_1$  and  $\underline{b}_1$  to denote the parameters  $\underline{\mathbf{w}}$  and  $\underline{b}$  in  $h(\mathbf{x})$  for the Fisher discriminant, and use  $\underline{\mathbf{w}}$  and  $\underline{b}$  for the SVM.

*1) Fisher Discriminant Classifier:*

The Fisher discriminant is based on the principle of projecting the data onto a one-dimensional space so that the two classes are well separated [21]. The discriminant vector  $\underline{\mathbf{w}}_1$  in the decision function is determined by

$$\underline{\mathbf{w}}_1 = \underline{\Sigma}^{-1}(\underline{\boldsymbol{\mu}}_{+1} - \underline{\boldsymbol{\mu}}_{-1}), \quad (3)$$

where  $\underline{\boldsymbol{\mu}}_{+1}$  and  $\underline{\boldsymbol{\mu}}_{-1}$  are the mean vectors of the two classes, and  $\underline{\Sigma}$  is the total within-class covariance matrix, all estimated from the training samples. The constant  $b_1$  is computed as

$$b_1 = -\frac{1}{2} \underline{\mathbf{w}}_1^T (\underline{\boldsymbol{\mu}}_{+1} + \underline{\boldsymbol{\mu}}_{-1}). \quad (4)$$

*2) SVM Classifier:*

SVM is a constructive learning procedure based on statistical learning theory [22]. It is based on the principle of structural risk minimization, which aims at minimizing the bound on the generalization error (i.e., error made by the learning machine on data unseen during training) rather than minimizing the mean square error over the data set. As a result, an SVM tends to perform well when applied to data outside the training set. In recent years, SVM learning has found a wide range of real-world applications (see, for example, [23-27]). In many of these applications it has been reported that SVM based approaches are able to outperform competing methods. In our own work [28] we developed an SVM based approach for detection of microcalcifications in mammograms, and demonstrated using clinical mammogram data that such an approach could outperform several well-known methods in the literature.

Using the training data set  $\{(\mathbf{x}_i, d_i), i = 1, 2, \dots, l\}$ , a linear SVM classifier in its original form is formulated as minimization of the following cost function:

$$J(\mathbf{w}, \xi) = \frac{1}{2} \mathbf{w}^T \mathbf{w} + C \sum_{i=1}^l \xi_i, \quad (5)$$

$$\text{subject to } d_i(\mathbf{w}^T \mathbf{x}_i + b) \geq 1 - \xi_i, \xi_i \geq 0; \quad i = 1, 2, \dots, l,$$

where  $C$  is a user-specified, positive parameter, and  $\xi_i$  are slack variables.

The cost function in (5) constitutes the so-called *structured risk*. It consists of both the empirical risk (i.e., the training errors reflected by the second term) and the model complexity measure (the first term). The regularization parameter  $C$  in (5) is used to define the trade-off between these two factors. In particular, when the two classes are separable, the SVM classifier amounts to maximize the separating margin between the two classes (as illustrated in Fig. 3(a)).

For our task at hand, we propose to modify the SVM cost function in (5) as:

$$J(\mathbf{w}, \xi) = \frac{1}{2} \mathbf{w}^T \mathbf{w} + C^+ \sum_{i \in Z^+} \xi_i + C^- \sum_{i \in Z^-} \xi_i, \quad (6)$$

where  $C^+ > C^-$ , and  $Z^+, Z^-$  are the index sets of the training samples belonging to class 1 (i.e.,  $d_i = +1$ ) and class 2 (i.e.,  $d_i = -1$ ), respectively. This imposes a greater penalty ( $C^+$ ) on missed similar images than on misclassified non-similar images ( $C^-$ ). The rationale is that the first-stage classifier is for pre-screening only and should be designed to pass marginal cases to the second stage for further consideration.

Using the technique of Lagrange multipliers, one can show that a necessary condition for minimizing  $J(\mathbf{w}, \xi)$  in (6) is that the vector  $\mathbf{w}$  is formed by a linear combination of the vectors  $\mathbf{x}_i$ , i.e.,

$$\mathbf{w} = \sum_{i=1}^l \alpha_i y_i \mathbf{x}_i, \quad (7)$$

where  $\alpha_i \geq 0, i = 1, 2, \dots, l$ , are the Lagrange multipliers associated with the constraints in (5).

The Lagrange multipliers  $\alpha_i \geq 0, i = 1, 2, \dots, l$ , are solved from the dual form of (6), which is expressed as

$$\max \tilde{J}(\alpha_1, \alpha_2, \dots, \alpha_l) = \sum_{i=1}^l \alpha_i - \frac{1}{2} \sum_{i=1}^l \sum_{j=1}^l \alpha_i \alpha_j d_i d_j \mathbf{x}_i^T \mathbf{x}_j \quad (8)$$

subject to the constraints:

$$\begin{aligned} (1) \quad & \sum_{i=1}^l \alpha_i d_i = 0 \\ (2) \quad & 0 \leq \alpha_i \leq C^+ \text{ for } d_i = +1; \text{ or } 0 \leq \alpha_i \leq C^- \text{ for } d_i = -1, \quad i = 1, 2, \dots, l. \end{aligned} \quad (9)$$

The cost function  $\tilde{J}(\alpha_1, \alpha_2, \dots, \alpha_l)$  is convex and quadratic in terms of the unknown parameters  $\alpha_i$ . In practice, the maximization in (8) is solved numerically through quadratic programming [22].

Analytic solutions of (8) are not readily available, but it is still informative to examine the conditions under which an optimal solution is achieved. The Karush-Kuhn-Tucker (KKT) optimality conditions for (8) lead to the following three cases for each  $\alpha_i$ :

1.  $\alpha_i = 0$ . This corresponds to  $d_i(\mathbf{w}^T \mathbf{x}_i + b) > 1$ . In this case, the data element  $\mathbf{x}_i$  is outside the decision margin of the function  $h(\mathbf{x})$  and is correctly classified.
2.  $0 < \alpha_i < C^+$  for  $d_i = +1$ ; or  $0 < \alpha_i < C^-$  for  $d_i = -1$ . In this case,  $\mathbf{w}^T \mathbf{x}_i + b = d_i$ . The data element  $\mathbf{x}_i$  is strictly located on the decision margin of  $h(\mathbf{x})$ . Hence,  $\mathbf{x}_i$  is called a *margin support vector* of  $h(\mathbf{x})$ .
3.  $\alpha_i = C^+$  for  $d_i = +1$ ; or  $\alpha_i = C^-$  for  $d_i = -1$ . In this case,  $d_i(\mathbf{w}^T \mathbf{x}_i + b) < 1$ . The data element  $\mathbf{x}_i$  is inside the decision margin (though it may still be correctly classified). Accordingly,  $\mathbf{x}_i$  is called an *error support vector* of  $h(\mathbf{x})$ .

It is typical that most of the training examples are correctly classified by the trained classifier (case 1), i.e., only a few training examples will be support vectors. For simplicity, let  $\mathbf{s}_j$ ,  $\alpha_j^*$ ,  $j = 1, 2, \dots, l_s$ , denote these support vectors and their corresponding nonzero Lagrange multipliers, respectively, and let  $d_j$  denote their class labels. The SVM decision function can thus be simplified as

$$h(\mathbf{x}) = \sum_{j=1}^{l_s} (\alpha_j^* y_j) \mathbf{s}_j^T \mathbf{x} + b. \quad (10)$$

Note that the decision function is now determined directly by the support vectors  $\mathbf{s}_j, j=1,2,\dots,l_s$ , which are determined by solving the optimization problem in (8) during the training phase.

### B. Regression Stage

The regression stage is used to provide quantitative *SCs* between the query and those images deemed sufficiently similar by the classification stage. Consequently, only a subset of the training data will be qualified for the training of the learning machine in this stage. In this study we consider the following two different approaches for learning the similarity function  $f(\mathbf{x})$ : (1) an SVM, and (2) a general regression neural network (GRNN) [29].

#### 1) SVM Regression:

SVM learning can also be applied for regression. An SVM formulation in such a case maintains many of the characteristics of the classification case. For nonlinear regression, an SVM in concept first maps the input data vector  $\mathbf{x}$  into a higher dimensional space  $\mathcal{H}$  through an underlying nonlinear mapping  $\Phi(\cdot)$ ; then applies a linear regression in this mapped space. That is, a nonlinear SVM regression function can be written in the following form:

$$f(\mathbf{x}) = \mathbf{w}^T \Phi(\mathbf{x}) + b. \quad (11)$$

Let  $\{(\mathbf{x}_i, y_i), i=1,2,\dots,l'\}$  denote a set of training samples surviving the first stage, where  $y_i$  is the human-observer *SC* for the image pair denoted by  $\mathbf{x}_i$ . The parameters  $\mathbf{w}$  and  $b$  in the regression function in (11) are determined through minimization of the following structured risk:

$$R(\mathbf{w}, b) = \frac{1}{2} \mathbf{w}^T \mathbf{w} + C \sum_{i=1}^{l'} L_\varepsilon(\mathbf{x}_i), \quad (12)$$

where  $L_\varepsilon(\cdot)$  is the so-called  $\varepsilon$ -insensitive loss function which is defined as:

$$L_\varepsilon(\mathbf{x}) = \begin{cases} |y - f(\mathbf{x})| - \varepsilon, & \text{if } |y - f(\mathbf{x})| \geq \varepsilon \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

The function  $L_\varepsilon(\cdot)$  has the property that it does not penalize errors below the parameter  $\varepsilon$ , as illustrated in Fig. 3(b). The constant  $C$  in (12) determines the trade-off between the model complexity and the training error.

As with the case of classification, the regression function  $f(\mathbf{x})$  in (11) is also characterized by the support vectors. It can be written as follows

$$f(\mathbf{x}) = \sum_{i=1}^{l'_s} \gamma_i K(\mathbf{x}, \mathbf{s}_i) + b, \quad (14)$$

where  $\mathbf{s}_i, i = 1, 2, \dots, l'_s$ , denote the support vectors, and  $K(\mathbf{x}, \mathbf{s}_i) \equiv \Phi(\mathbf{x})^T \Phi(\mathbf{s}_i)$  which is called a kernel function. A training sample  $(\mathbf{x}_i, y_i)$  is a *margin support vector* when  $|f(\mathbf{x}_i) - y_i| = \varepsilon$ , and an *error support vector* when  $|f(\mathbf{x}_i) - y_i| > \varepsilon$ .

From (14), we can directly evaluate the regression function through the kernel function  $K(\cdot, \cdot)$  without the need to specifically addressing the underlying mapping  $\Phi(\cdot)$ . In this study we consider two kernel types: polynomial kernels and Gaussian radial basis functions (RBF). These are among the most commonly used kernels in SVM research, and are known to satisfy Mercer's condition [22]. They are defined as follows.

1. Polynomial kernel:

$$K(\mathbf{x}, \mathbf{y}) = (\mathbf{x}^T \mathbf{y} + 1)^p, \quad (15)$$

where  $p > 0$  is a constant that defines the kernel order.

2. RBF kernel:

$$K(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{\|\mathbf{x} - \mathbf{y}\|^2}{2\sigma^2}\right), \quad (16)$$

where  $\sigma > 0$  is a constant that defines the kernel width.

2) *GRNN Regression:*

The GRNN computes an estimate of the conditional mean of the  $SC$  for an image pair from the human-observer data [29]. It is based on an estimate of the joint probability density of the input and the output obtained by the Parzen method [29]. With training data  $\{(\mathbf{x}_i, y_i), i = 1, 2, \dots, l'\}$ , the output of the GRNN can be represented as:

$$f(\mathbf{x}) = \frac{\sum_{i=1}^{l'} y_i \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}_i\|^2}{2\sigma^2}\right)}{\sum_{i=1}^{l'} \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}_i\|^2}{2\sigma^2}\right)}, \quad (17)$$

where  $\sigma > 0$  defines the kernel width.

Note that the GRNN estimate  $f(\mathbf{x})$  in (17) has a similar form to the SVM estimate in (14) when the RBF kernel is used. The major difference between the two, however, is that only the support vector samples are used in the SVM in (14), while all the training samples are used in the GRNN in (17). Thus, the SVM estimate can be computationally advantageous over the GRNN.

#### IV. RELEVANCE FEEDBACK

In this section we explore how to incorporate relevance feedback into our proposed learning based retrieval approach. Relevance feedback is a post-query process to refine the search by using positive and/or negative indications from the user of the relevance of retrieved images. It has been applied successfully in traditional text-retrieval systems for improving the results of a retrieval strategy [30]. In particular, we consider the following scenario: for a query image  $\mathbf{q}$  a user selects a relevant image  $\mathbf{r}$  amongst the retrieved images to confirm that the retrieved  $\mathbf{r}$  is indeed similar to the query  $\mathbf{q}$ ; we want to incorporate this information to further refine the search, hoping that more relevant images could be found for the same query  $\mathbf{q}$ .

In this study we consider the following simple approach for relevance feedback: we explicitly weigh in the impact of the feedback image  $\mathbf{r}$  in the similarity between the query image  $\mathbf{q}$  and a database entry  $\mathbf{d}$ . Specifically, we use the following weighted  $SC$ :

$$\widetilde{SC}(\mathbf{q}, \mathbf{d}) = (1 - w) \cdot SC(\mathbf{q}, \mathbf{d}) + w \cdot SC(\mathbf{r}, \mathbf{d}), \quad (18)$$

where  $w$  is a weighting parameter used to adjust the relative impact of the feedback image  $\mathbf{r}$ . The images with the highest weighted  $SC$ s are then retrieved.

An alternative to the above weighting approach is to adapt the learning machine based on the feedback information. We will consider this in a separate study [31], as the main goal of this paper is to demonstrate the feasibility of a learning framework for similarity modeling.

## V. PERFORMANCE EVALUATION STUDY

### A. Mammogram Data Set

The proposed retrieval framework was developed and tested using a data set collected by the Department of Radiology at The University of Chicago. This data set consists of 76 clinical mammograms, all containing multiple microcalcifications (MCs). These mammograms are of dimension  $1000 \times 700$  pixels, with a spatial resolution of 0.1 mm/pixel and 10-bit grayscale. Collectively, there are a total of 1120 MCs in these mammograms, which were identified by a group of experienced mammographers.

MCs are tiny calcium deposits that appear as small bright spots (typically 0.05-1mm in diameter) in a mammogram. MC clusters (MCCs) in a mammogram provide valuable information to radiologists in diagnosis of cancer. For example, linearly distributed MCCs are typically malignant, while round clusters are typically benign [32]. In Fig. 4 we show a number of different regions of interests (ROIs) extracted from the mammograms in the data set, all of which contain MCCs.

Our objective is to apply the proposed framework to retrieve mammograms containing similar MCCs to that in a query mammogram.

### B. Observer Similarity Data

For the training and testing of the algorithms, ROIs containing the identified MCCs were first extracted from all the mammograms in the data set (as shown in Fig. 4). Among the 76 mammograms, 74 contain only single ROIs, while the other two have two ROIs. These MCC ROIs were then used in a

subsequent observer study to obtain SCs for different ROI pairs, which were then used to form training and testing samples.

The observer study was carried out by a panel of six human observers, who scored the similarity between each pair of ROIs based on their geometric distributions on a scale from 0 (most dissimilar) to 10 (most similar). It consisted of the following different sessions: 1) a “pre-calibration” session; 2) individual scoring sessions; and 3) a statistical analysis session for both intra- and inter-observer consistencies.

The panel of observers first participated in a “pre-calibration” session (~1 hour), the goal of which was to establish a consensus among the observers on a uniform measure of the perceptual similarity and to identify tentative “anchor pairs” (prototype examples) along the scale (from very different to very similar, all chosen randomly from the mammogram set).

For the individual scoring sessions, we randomly selected 30 ROIs from the mammogram set, each of which corresponds to a different patient. The observers then scored the similarity for all the possible pairs (a total of 435) formed by these ROIs, assisted by a software user interface. In each session, a query ROI was displayed along with up to 15 other ROIs simultaneously on the same computer screen (presented in a random order). The observer then assigned a continuous SC value between the query and each of the other ROIs by mouse-clicking on a thermometer bar on the computer screen. Each of the 30 ROIs in the dataset was used in turn as the query, yielding a total of 870 SCs (each MCC pair got scored twice) from each observer.

The collected SCs were then analyzed for both intra- and inter-observer consistencies. Specifically, the Kendall’s rank correlation method [33] was first applied to test the consistency between the two scores for each of the 435 MCC pairs by the same observer. The two scores were then averaged for each pair. Afterward, the Kendall’s rank correlation method was applied to analyze the inter-observer consistency among the six observers. The scores by the six observers were then averaged for each of the 435 pairs to obtain the SCs. ([Issam, describe all the test results here.](#))

Finally, we introduced 30 “ideal” pairs, formed by each query ROI with itself. These pairs were all assigned a perfect SC (10).

In summary, a total of 465 MCC pairs were scored and recorded. In Fig. 5 we show a histogram plot of the obtained observers' similarity scores for the 465 pairs.

All the six observers have background in engineering training. To facilitate the observer study, all the individual MCs were clearly marked out in all the ROIs involved. While there are other potentially important image features one might consider, we elected in this preliminary demonstration to retrieve images based on the spatial characteristics of the clusters alone. This will enable us to demonstrate the feasibility of the proposed framework using an observer-data set of reasonable size.

### C. Extraction of MCC Features

To describe the geometric features of MCCs, we started with a pool of 20+ (Issam, give your specific number) shape descriptors, most of which were used in the literature for shape analysis of MMCs [xx-yy]. We then applied a so-called *sequential backward selection* procedure [34] to reduce the set of salient features down to nine, which we describe below in detail. We point out that features xx-yy were used in [xx], features were used in [xx]. (Issam, update the references above)

1. *Cross sectional area (A)*: the area occupied by the cluster. It is computed in the following steps: 1) a binary image is first created in which the pixels corresponding to the centers of the MCs are set to 1 and all the rest of the pixels set to 0. 2) a Delaunay triangulation is next applied to connect the centers of the MCs in this binary image; the average inter-distance between neighboring MCs, denoted by  $\bar{\rho}$ , is then computed based on this triangulation. 3) a morphological closing operation with a circular structuring element having a radius of  $\bar{\rho}$  is then performed on the binary image to fill the gaps among the MCs. The area of the resulting region is then computed.
2. *Compactness*: a measure of roundness of the region occupied by the cluster. It is computed as

$$C_f = \frac{4\pi A_f}{P_f^2}, \quad (21)$$

$A_f, P_f$  are the area and perimeter of the solid region occupied by the cluster (i.e., holes are filled when necessary), respectively. Note that  $A_f$  will differ from the cross sectional area  $A$  when the occupied region contains any holes.

3. *Eccentricity*: the eccentricity of the smallest enclosing ellipse of the region, computed as the ratio of the distance between the foci and the length of the major axis of the ellipse.
4. *Density*: the spatial density of the MCs in the cluster, computed as the number of MCs per unit area ( $A$ ).
5. *Scatteredness*: represented by the mean and the standard deviation of the inter-distances between neighboring MCs; the neighbors are determined based on the Delaunay triangulation of the MCs as described above.
6. *Solidity*: computed as the ratio between cross sectional area  $A$  and the area of the convex hull formed by the MCs.
7. *Invariant moment  $\phi_1$* : a regional descriptor that is invariant to translation, rotation, or scaling [35].
8. *Moment signature*, as defined in [36]: a measure of boundary roughness, computed based on the distance deviation of a point on the boundary from the center of the region.
9. *Normalized Fourier descriptor*, also as defined in [36]: a frequency-domain characterization of the smoothness of the boundary.

These feature components (a total of 10, with 2 for scatteredness) were first computed for each MCC in the mammograms. All these feature components were then normalized to have the same dynamic range (0,1). Each MCC was then labeled with a feature vector  $\mathbf{u}$  formed by these components. These feature vectors were paired with the observer similarity data to form the training and testing samples.

In summary, we have the following data set

$$S = \{(\mathbf{x}_i, y_i), i = 1, 2, \dots, 465\}, \quad (22)$$

where  $\mathbf{x}_i$  denotes the computed feature vector for the  $i$ -th MCC pair, and  $y_i$  is the observer  $SC$  of the pair. This set was used for the subsequent training and testing of the proposed framework.

#### D. Machine Training and Performance Evaluation

##### 1) Preparation of Data Sets:

For the first stage, the MCC pairs in set  $S$  in (22) were first divided into two classes: class 1 representing sufficiently similar pairs, and class 2 representing dissimilar pairs. We chose a threshold  $T_1=4$  so that samples in  $S$  were labeled as class 1 if their SCs were larger than  $T_1$ ; otherwise they were labeled as class 2. In short, we denote this set as

$$S_1 = \{(\mathbf{x}_i, d_i), i = 1, 2, \dots, 465\}, \quad (23)$$

where  $d_i = \text{sgn}(y_i - T_1)$ . There were in total 229 samples in class 1, and 236 samples in class 2. This set was used subsequently to train and test the first-stage classifier.

For the regression stage, we chose only those pairs in  $S$  with SCs larger than  $T_1$ , i.e., those belonging to class 1 in set  $S_1$ . We denote this set as

$$S_2 = \{(\mathbf{x}_i, y_i) : y_i > T_1, i = 1, 2, \dots, 465\}. \quad (24)$$

##### 2) Performance Evaluation:

For training and testing of the learning machines (both the classification stage and the regression stage), we applied the following cross-validation procedure [37]: 1) the images were selected in turn so that during each run only one image was chosen (as a query), based on which the data samples ( $S_1$  or  $S_2$ ) were divided into the following two sets: one for training, which consisted of all the samples not involving the chosen image, and the other for testing, which consisted of only those samples involving the chosen image; 2) in each run the learning machine (either classification or regression) was then trained using the resulting training set, and tested for performance using the testing set; 3) the test results were then averaged over all the different runs to obtain the generalization performance (e.g., classification error, retrieval precision, etc.).

To evaluate the performance of the retrieval network, we used the so-called *precision-recall curves* [1]. The retrieval *precision* is defined as the proportion of the images among all the retrieved that are truly

relevant to a given query; the term *recall* is measured by the proportion of the images that are actually retrieved among all the relevant images to a query. The precision-recall curve is a plot of the retrieval precision vs the recall over a continuum of the operating threshold  $T$  (Fig. 1).

As the ground truth in calculation of the precision-recall curves, we considered an image to be truly relevant to a query provided that their corresponding observer  $SC$  is larger than a pre-selected threshold  $T_2$ . In our experiments,  $T_2 = 6$  was used.

### 3) *Relevance Feedback:*

To demonstrate the effect of relevance feedback, we performed the following experiments: for each query, the trained retrieval network was first applied to retrieve images from the database; among the images retrieved, the one with the highest  $SC$  (based on the pre-existing observer data) was chosen as the relevant feedback image (in case there was a tie, random selection was used to break the tie). This chosen image was then paired up with the query along with their corresponding observer  $SC$  to form the feedback sample. The proposed relevance feedback procedure was then applied to retrieve a new set of images. The precision-recall curves were then computed based on this new set of images.

### 4) *Impact of parameters:*

To demonstrate the impact of various parameters involved in the training and testing of the proposed network on the overall performance, we also evaluated the precision-recall curves when these parameters were varied. This includes using different values for the thresholds  $T_1$  and  $T_2$ , and perturbing the inherent parameters of the trained learning machine.

## VI. EXPERIMENTAL RESULTS

The proposed two-stage learning approach was thoroughly tested and evaluated for retrieval under various learning-machine settings. We summarize the results in Fig. 6 using the precision-recall curves for the following different network structures:

- 1) a linear Fisher discriminant for the first stage and an SVM for the second stage (Fisher-SVM);
- 2) a linear SVM for the first stage and an SVM for the second stage (SVM-SVM);

- 3) a linear SVM with the modified objective function in (6) for the first stage and an SVM for the second stage (MSVM-SVM);
- 4) a linear SVM with the modified objective function in (6) for the first stage and a GRNN for the second stage (MSVM-GRNN).

In the first three structures a Gaussian kernel was used in the SVM for the second stage. We note that similar performance was also achieved when a polynomial kernel was used, of which the results are omitted for clarity of the plots.

For comparison, we also show in Fig. 6 the precision-recall curve obtained when a single stage SVM regression network was used for retrieval (SVM). In this case, the SVM with a Gaussian kernel was trained and tested directly using the samples formed from the entire set of observer SCs.

Moreover, we show in Fig. 6 the precision-recall curve obtained when a naïve Euclidian metric was used as the similarity measure. In this case, the images with features vectors closest to a query were retrieved.

From these results we see that the two-stage network (MSVM-SVM) achieves the best performance; and all the learning based networks outperform that based on the Euclidian distance. Note that the precision-recall curves corresponding to both Fisher-SVM and SVM-SVM drop below that of the single-stage network (SVM) as the recall ratio is increased toward unity. This can be explained as follows: at a fixed operating threshold the first stage Fisher or SVM classifier discards some of the relevant images for a query with a nonzero probability, preventing the recall ratio from reaching 1 (as the retrieval threshold  $T$  gets decreased). The use of a modified SVM classifier in the first stage avoids this pitfall.

Finally, the parametric settings for the learning machines corresponding to each of the network structures above are listed in Table 1. In our experiments each network structure was studied over a wide range of parametric settings; the precision-recall curves in Fig. 6 represent the best results for each case. The performance of the learning networks was found to be considerably robust and insensitive to changes in the parameters (such as regularization parameter  $C$ , insensitive bound  $\varepsilon$ , kernel width  $\sigma$ , etc.). As an

example, in Fig. 7 we show the resulting precision-recall curves for the case of MSVM-SVM when the parameter  $C$  of the regression stage is varied from its listed value in Table 1.

**Table 1.** PARAMETRIC SETTINGS OF THE TRAINED RETRIEVAL NETWORKS

Networks	Classification Stage	Regression Stage
Fisher-SVM	-	$C=100, \sigma=1.5, \varepsilon=0.5$
SVM-SVM	$C=1000$	$C=100, \sigma=1.5, \varepsilon=0.5$
MSVM-SVM	$C^+=10^5, C^-=5 \times 10^4$	$C=100, \sigma=1.5, \varepsilon=0.5$
MSVM-GRNN	$C^+=10^5, C^-=5 \times 10^4$	$\sigma=0.25$
SVM	-	$C=100, \sigma=1.5, \varepsilon=0.5$

In Figs. 8 and 9 we show some retrieval examples for two given query images. These results demonstrate that the two-stage network can indeed improve retrieval performance.

In Fig. 10, we show the precision-recall curves obtained using the Weighted  $SC$  method for the two-stage MSVM-GRNN. Similar results were also obtained for SVM and MSVM-SVM (but not shown for brevity). As can be seen, the proposed feedback procedures can further improve the retrieval performance.

In our experiments the two-stage networks (MSVM-SVM and MSVM-GRNN) could provide 4-5 times of saving in computation time in retrieval over that of a single-stage network (SVM).

In addition, we show in Fig. 11 the precision-recall curves obtained for MSVM-SVM when different values for the thresholds  $T_1$  and  $T_2$  were used.

Finally, in Fig. 12 we show the average percentage of images among retrieved that actually match the disease condition of the query, obtained using the leave-one-out procedure, when different values were used for the operating threshold  $T$ . For comparison, we also show in Fig. 12 correspondingly the matching percentage when the observer score (ground truth) is used. As can be seen, the developed retrieval system matches well with the observer, and, as expected, the matching percentage decreases as the threshold  $T$  decreases.

## VII. CONCLUSIONS

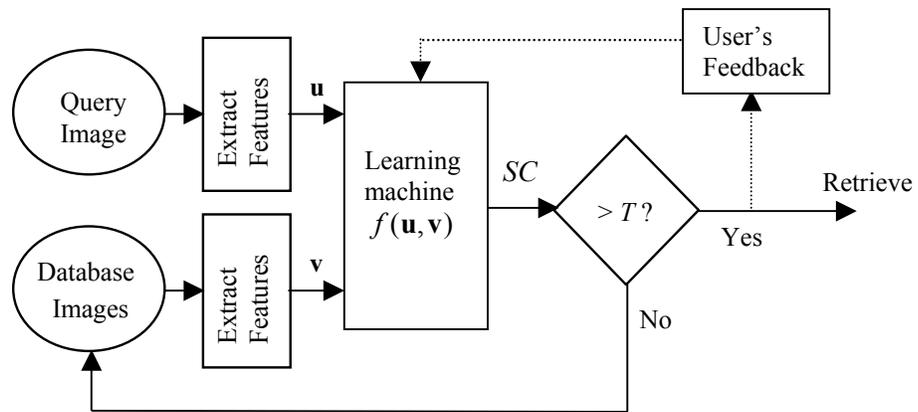
In this paper we have proposed a learning machine-based framework for modeling human perceptual similarity for content-based image retrieval. The proposed approach was developed and evaluated for retrieval of clinical mammograms containing clustered microcalcifications. The results demonstrated that a learning framework can be used effectively to model the perceptual similarity, thereby serving as basis for retrieving visually similar mammograms from a database. It was demonstrated that a hierarchical two-stage learning network can offer several advantages over a single-stage one, including faster speed and retrieval accuracy. Furthermore, the use of relevance feedback in such a framework can be used to further improve the retrieval performance. In our future work we will explore the use of incremental learning to adapt the learning network online to a user's feedback; we will also investigate the clinical benefit of using the developed retrieval framework for computer-aided diagnosis.

## REFERENCES

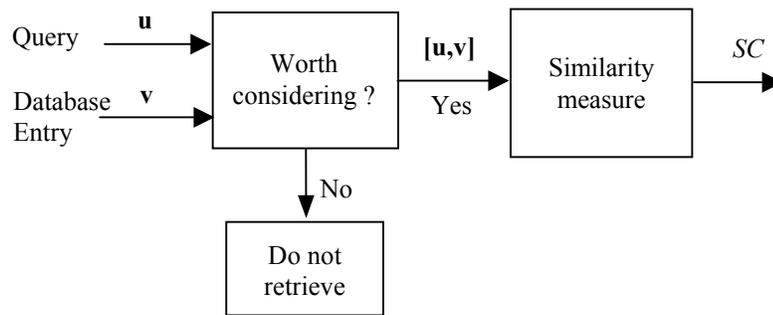
- [1] A. D. Bimbo, *Visual Information Retrieval*, Morgan Kaufman Publishers, 1999.
- [2] Y. Rui and T. Huang, "Image retrieval: current techniques, promising directions and open issues," *Journal of Visual Communication and Image Representation*, vol. 10, pp. 39-62, 1999.
- [3] T. Kato, T. Kurita, N. Otsu, and K. Hirata, "A sketch retrieval method for full color image database," in *Proceedings of 11<sup>th</sup> International Conference on Pattern Recognition*, The Hague, pp. 530-533, 1992.
- [4] W. Niblack, R. Barber, W. Equitz, M. Flickner, E. Glasman, D. Petkovic, P. Yanker, and C. Faloutsos, "The QBIC project: Querying images by content using color, texture, and shape," in *SPIE Proceeding on Storage and Retrieval for Image and Video Databases*, vol. 1908, pp. 137-46, 1993.
- [5] S. Wong, "CBIR in medicine: still a long way to go," *IEEE workshop on Content-Based Access of Image and Video Libraries*, Santa Barbara, CA, pp. 115-115, June 1998.
- [6] H. A. Swett and P. L. Miller, "ICON: a computer-based approach to differential diagnosis in radiology," *Radiology*, vol. 163, pp. 555-558, 1987.
- [7] P. M. Kelly and T. M. Cannon, "CANDID: comparison algorithm for navigating digital image databases," in *Proc. of the International Working Conference on Scientific and Statistical Database Management*, pp. 252-258, 1994.
- [8] A. Guimond and G. Subsol, "Automatic MRI database exploration and applications," *Pattern Recognition and Artificial Intelligence*, vol. 11, no. 8, pp. 1345-1365, 1997.

- [9] Y. Liu and F. Dellaert, "A classification based similarity metric for 3D image retrieval," in *Proc. of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 800-805, 1998.
- [10] J. Sklansky, E. Tao, M. Bazargan, C. Ornes, R. Murchison, and S. Teklehaimanot, "Computer-aided, case-based diagnosis of mammographic regions of interest containing microcalcifications," *Academic Radiology*, vol. 7, no. 6, pp. 395-406, 2000.
- [11] C. Ornes and J. Sklansky, "A visual neural classifier," *IEEE Trans. Systems, Man, and Cybernetics*, Part B, vol. 28, pp. 620-625, 1998.
- [12] Y. Kawata, N. Niki, H. Ohmatsu, M. Kusumoto, R. Kakinuma, K. Mori, H. Nishiyama, K. Eguchi, M. Kaneko, and N. Moriyama, "Three-dimensional CT image retrieval in a database of pulmonary nodules," in *Proc. of IEEE International Conference on Image Processing*, pp. 149-152, Rochester, NY, September, 2002.
- [13] I. El-Naqa, M. N. Wernick, Y. Yang, and N. P. Galatsanos, "Image retrieval based on similarity learning," in *Proceedings of IEEE International Conference on Image Processing*, pp. 722-725, Vancouver, Canada, 2000.
- [14] A. I. Mushlin, R. W. Kouides, and D. E. Shapiro, "Estimating the accuracy of screening mammography: a meta-analysis," *American Journal of Preventive Medicine*, vol. 14, no. 2, pp. 143-53, 1998.
- [15] R. N. Strickland and H.L. Hahn, "Wavelet transforms for detecting microcalcifications in mammograms," *IEEE Trans. on Medical Imaging*, vol. 15, no. 2, pp. 218-229, 1996.
- [16] E. A. Sickles, "Mammographic features of 300 consecutive non-palpable breast cancers," *American Journal of Roentgenology*, vol. 146, pp. 661-663, 1986.
- [17] D. B. Kopans, "The positive predictive value of mammography," *American Journal of Roentgenology*, vol. 158, pp. 521-526, 1992.
- [18] J. G. Elmore, C. K. Wells, C. H. Lee, D. H. Howard, and A. R. Feinstein, "Variability in radiologists' interpretations of mammograms," *New England Journal of Medicine*, vol. 331, no. 22, pp. 1493-1499, 1994.
- [19] R. M. Nishikawa, M. L. Giger, K. Doi, C. J. Vyborny, and R. A. Schmidt, "Computer aided detection of clustered microcalcifications in digital mammograms," *Medical and Biological Engineering and Computing*, vol. 33, pp. 174-178, 1995.
- [20] P. L. Miller, "Critiquing anesthetic management: the "ATTENDING" computer system," *Anesthesiology*, vol. 58, pp. 362-369, 1983.
- [21] B. Ripley, *Pattern Recognition and Neural Networks*, Cambridge, 1996.
- [22] V. Vapnik, *Statistical Learning Theory*, New York: Wiley, 1998.
- [23] M. Pontil and A. Verri, "Support vector machines for 3-D object recognition," *IEEE Trans. PAMI*, vol. 20, pp. 637-646, June 1998.
- [24] V. Wan and W. M. Campbell, "Support vector machines for speaker verification and identification," in *Proc. IEEE Workshop on Neural Networks for Signal Processing*, pp. 775-784, Sydney, Australia, December 2000.
- [25] E. Osuna, R. Freund, and F. Girosi, "Training support vector machines: application to face detection," in *Proc. Computer Vision and Pattern Recognition*, pp. 130-136, Puerto Rico, 1997.
- [26] T. Joachims, "Transductive inference for text classification using support vector machines," in *Proc. Int. Conf. on Machine Learning*, Slovenia, June 1999.

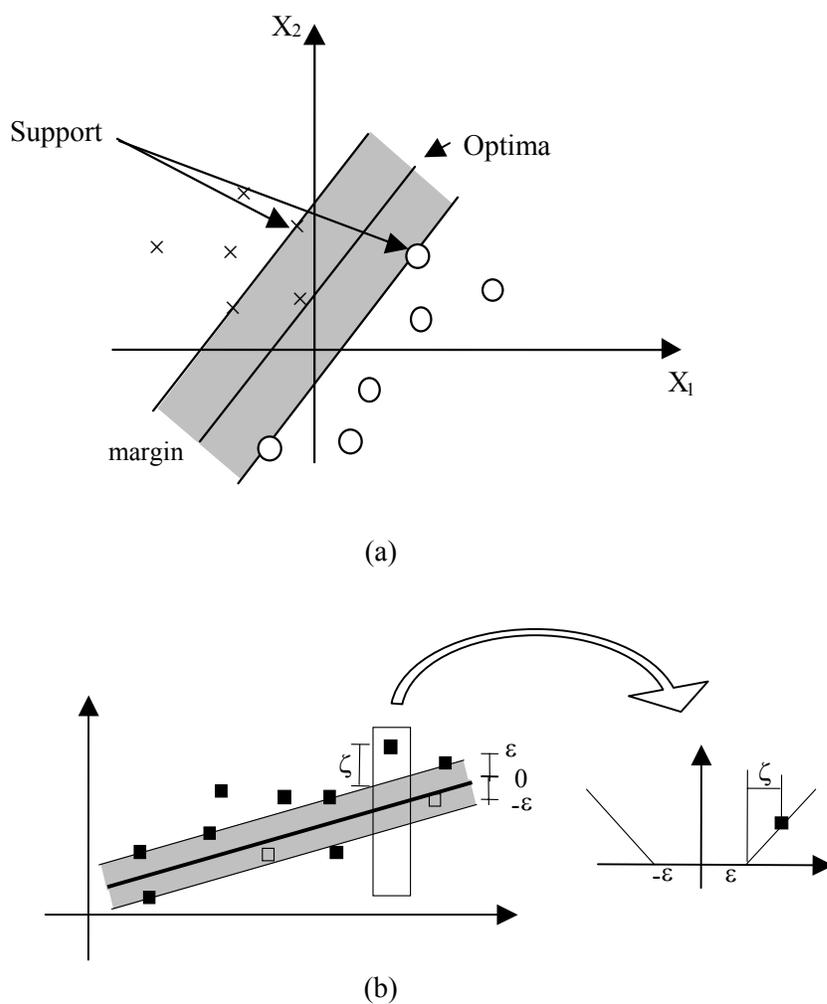
- [27] C. J. Burges, "A tutorial on support vector machines for pattern recognition," *Data Mining and Knowledge Discovery*, vol. 2, no. 2, 121-167, 1998.
- [28] I. El-Naqa, Y. Yang, M. N. Wernick, N. P. Galatsanos, "A support vector machine for approach for detection of Microcalcifications," *IEEE Trans. Med. Imaging*, vol. 21, no. 12, pp. 1552-1563, 2002.
- [29] D. F. Specht, "A general regression neural network," *IEEE Trans. Neural Network*, vol. 2, no. 16, pp. 568-576, 1991.
- [30] Y. Rui, T. Huang, M. Ortega, S. Mehrotra, "Relevance feedback: A power tool for interactive content-based image retrieval," *IEEE Trans. Circuits and Systems for Video Technology*, vol. 8, no. 5, Sept 1998.
- [31] I. El-Naqa, Y. Yang, N. P. Galatsanos, and M. N. Wernick, "Relevance feedback based on incremental learning for mammogram retrieval," *Proceedings of IEEE International Conference on Image Processing*, Barcelona, Spain, Sept. 13-17, 2003.
- [32] American College of Radiology, *Illustrated Breast Imaging Reporting and Data System*, 3<sup>rd</sup> ed., 1998.
- [33] M. G. Kendall, *Rank Correlation Methods*, London: Griffin, 4<sup>th</sup> Ed., 1970.
- [34] S. Theodoridis and K. Koutroumbad, *Pattern Recognition*, Academic Press, 2<sup>nd</sup> Ed., 2003.
- [35] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, Addison-Wesley, 1992.
- [36] S. Liang, R. Rangayyan, and J.E. Desautels, "Application of shape analysis to mammographic calcifications," *IEEE Trans. Med. Imaging*, vol. 13, no. 2, pp. 263-274, 1994.
- [37] R. Kennedy, Y. Lee, B. Van Roy, C. D. Reed, and R. P. Lippman, *Solving Data Mining Problems Through Pattern Recognition*, Prentice Hall, 1998.



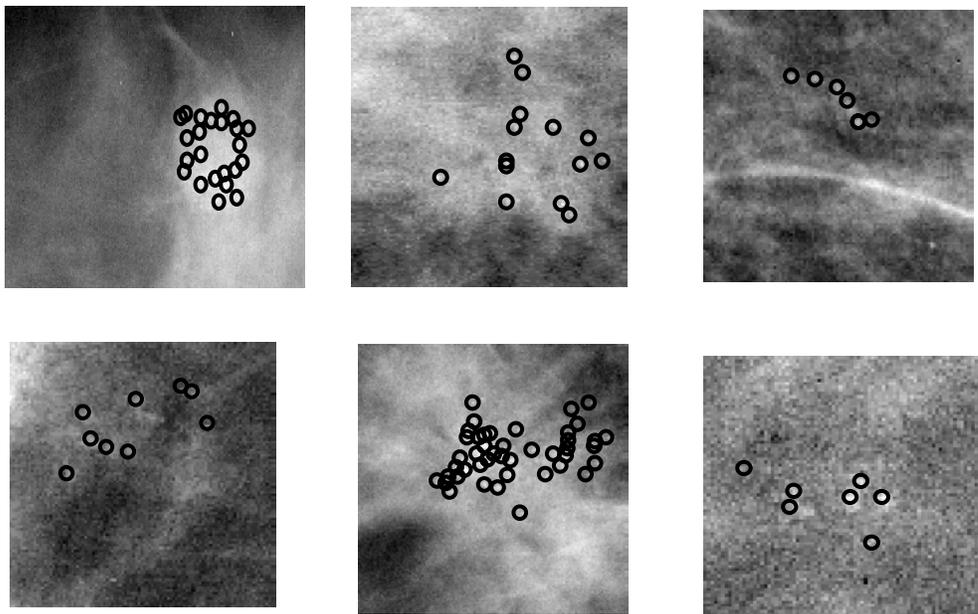
**Figure 1.** Proposed image retrieval framework with relevance feedback.



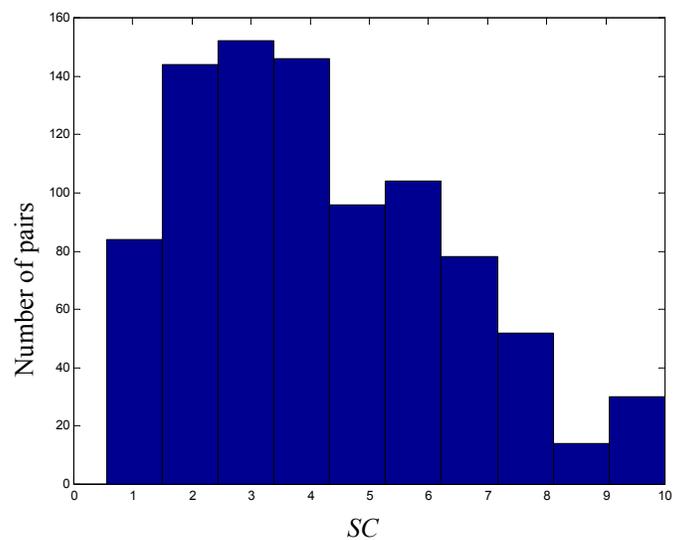
**Figure 2.** The two-stage hierarchical learning framework.



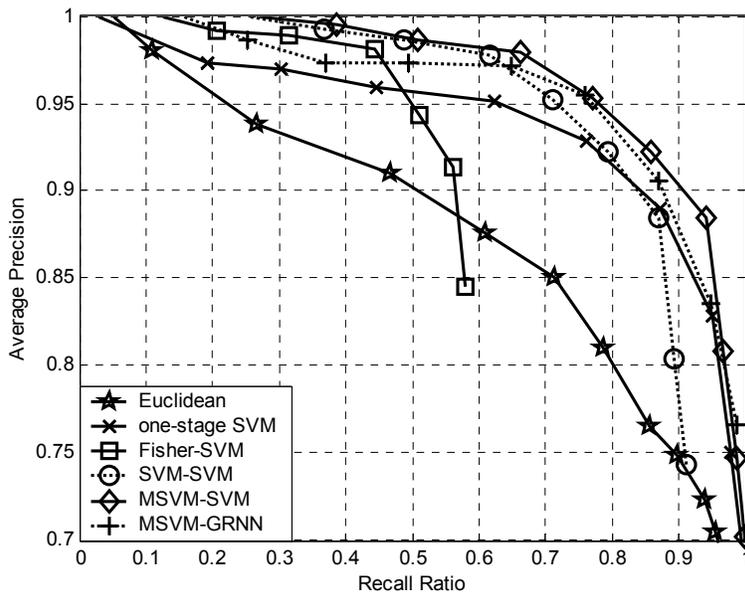
**Figure 3.** Illustration of support vector machines: (a) classification with a linear hyperplane that maximizes the margin between the two classes; and (b)  $\epsilon$ -insensitive SVM for regression, where the loss function does not penalize errors below the parameter  $\epsilon$ . The support vectors are indicated by filled squares.



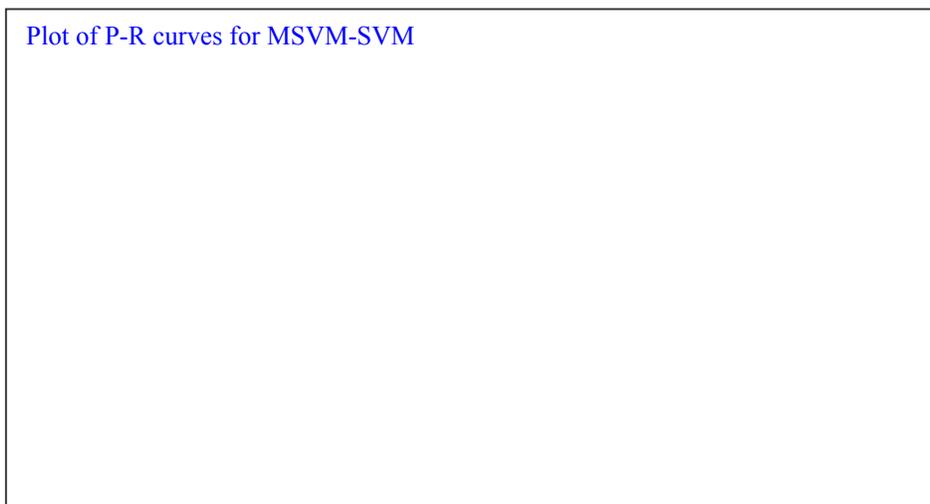
**Figure 4.** Examples of mammogram regions containing clustered microcalcifications (indicated by circles).



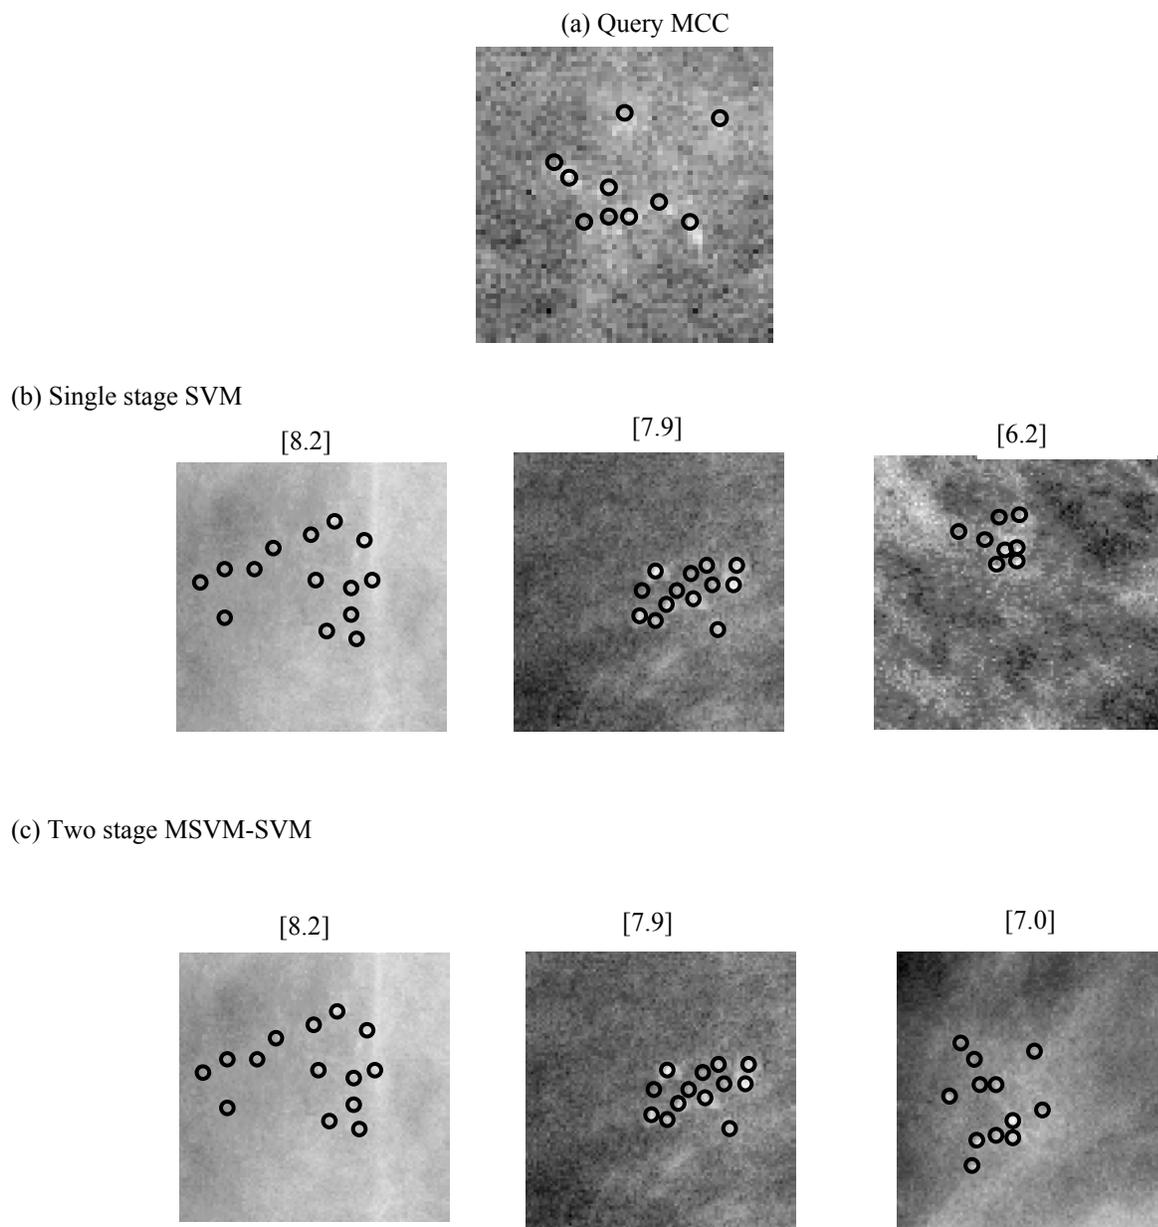
**Figure 5.** A histogram plot of the observers' similarity scores (SCs) for the 465 MCC pairs in the dataset.



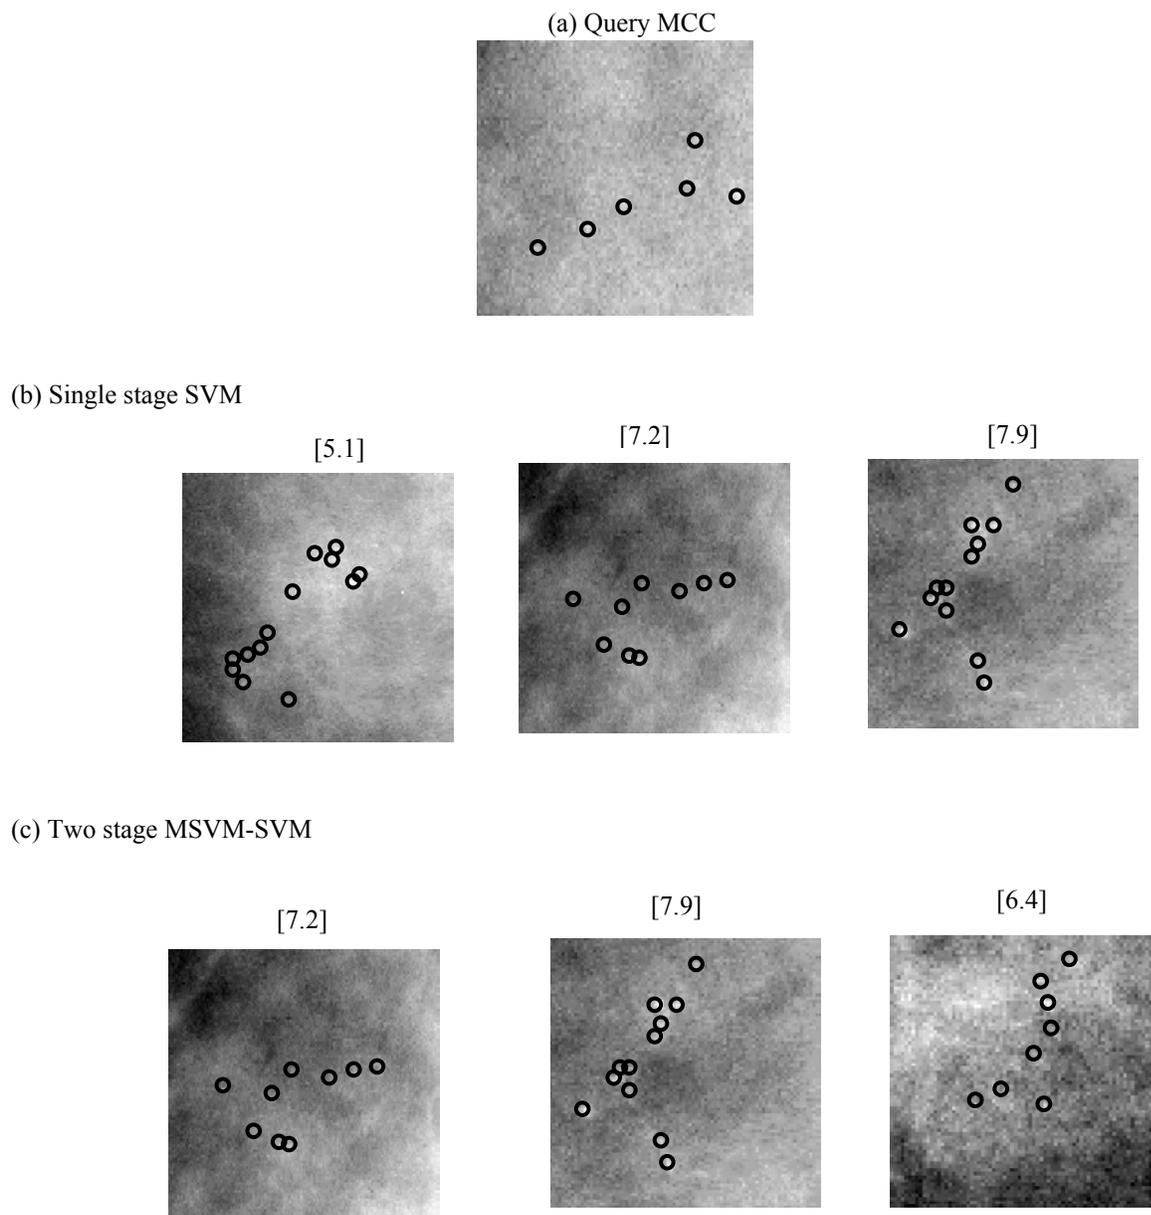
**Figure 6.** Plot of precision-recall curves obtained from various network structures.



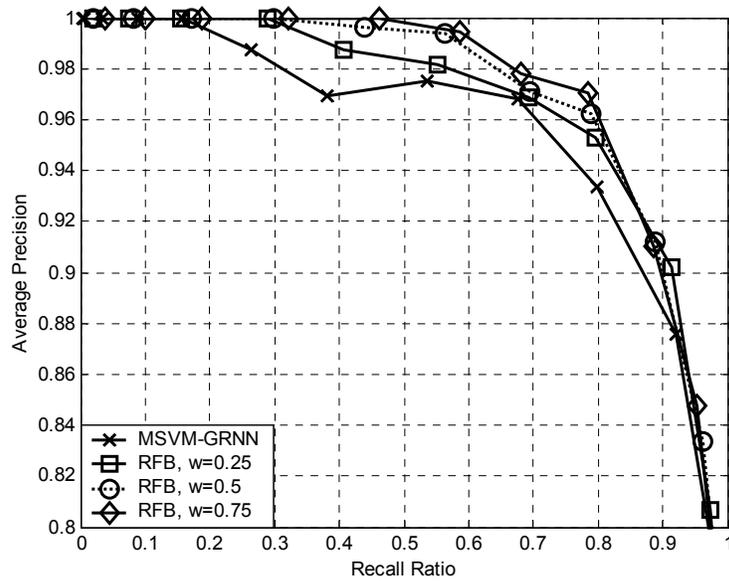
**Figure 7.** Precision-recall curves obtained by the MSVM-SVM network when the regularization parameter  $C$  in the regression stage is varied from its tuned value of 100 in Table 1 to: (a) 10 (MSVM-SVM(10)), and (b) 1000 (MSVM-SVM(1000)).



**Figure 8.** Top three images retrieved by the single-stage SVM (b) and the two-stage MSVM-SVM (c), respectively, for a given query MCC (a). Numbers in brackets on top of each cluster are the user SCs.



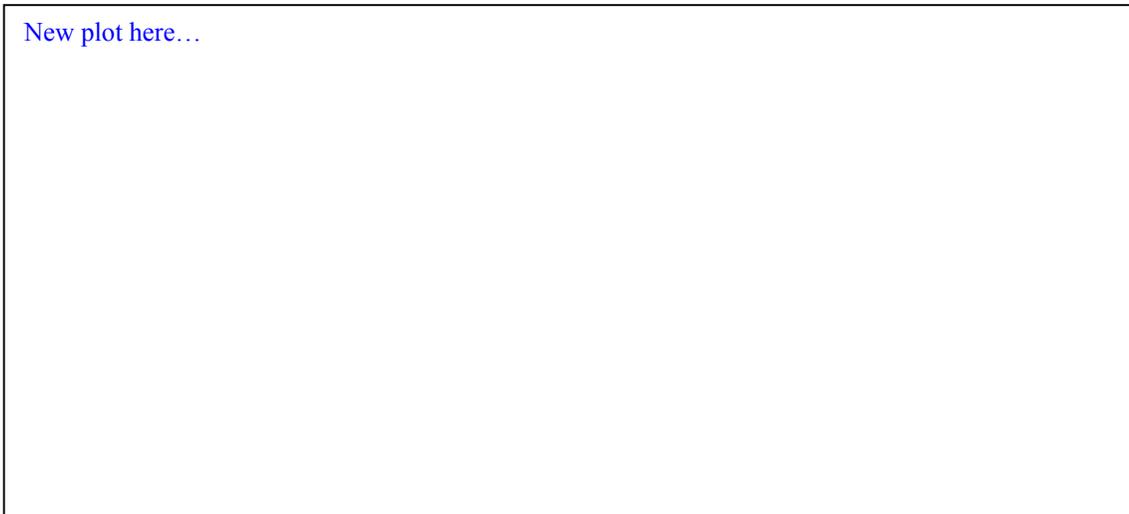
**Figure 9.** Top three images retrieved by the single-stage SVM (b) and the two-stage MSVM-SVM (c), respectively, for a given query MCC (a). Numbers in brackets on top of each cluster are the user SCs.



**Figure 10.** Precision-recall curves using relevance feedback for the two-stage network MSVM-GRNN.

New plot here...

**Figure 11.** Precision-recall curves obtained by the MSVM-SVM network when different values were used for the thresholds  $T_1$  and  $T_2$ : (a)  $T_1=3$ ,  $T_2=6$  (MSVM-SVM (3,6)); (b)  $T_1=4$ ,  $T_2=6$  (MSVM-SVM (4,6)); and (c)  $T_1=4$ ,  $T_2=7$  (MSVM-SVM (3,7)).



New plot here...

**Figure 12.** Plot of the average percentage of images among retrieved that actually match the disease condition of the query when different values were used for the operating threshold  $T$ . For comparison, we also show correspondingly the matching percentage when the observer score (ground truth) is used.