

Classification of Radiographs in the 'Image Retrieval in Medical Applications' - System (IRMA)

**J. Dahmen¹, T. Theiner¹, D. Keysers¹, H. Ney¹,
T. Lehmann², B. Wein³**

¹Lehrstuhl für Informatik VI

²Institute of Medical Informatics

³Department of Diagnostic Radiology

RWTH Aachen - University of Technology, Germany

Email: dahmen@informatik.rwth-aachen.de

Abstract

In this paper we present a new approach to classifying radiographs, which is the first important task of the IRMA system. Given an image, we compute posterior probabilities for each image class, as this information is needed for further IRMA processing. Classification is done by using an extended version of Simard's tangent distance within a kernel density based classifier. We propose a new distortion model for radiographs and prove its effectiveness by applying the method to 1617 radiographs coming from daily routine. Although the distortion model alone yields very good results, the best recognition rates are obtained by combining it with tangent distance, i.e. by computing a 'distorted tangent distance'.

1 Introduction

The importance of digital image retrieval techniques increases in the emerging fields of medical imaging and picture archiving and communication systems (PACS). Up to now, textual index entries are mandatory to retrieve medical images from a hospital archive, even if the archive is DICOM-compliant (Digital Imaging and Communications in Medicine [Lou et al., 1997]). Currently, a lot of research is done in the field of image retrieval, but the majority of today's systems are intended for browsing variegated image databases, e.g. for images collected from the World Wide Web. Usually, the features used for indexing characterize the entire image rather than image regions or objects and FALOUTSOS [Faloutsos et al., 1994] pointed out that one of the most effective features of such systems is color. Unfortunately, it is quite obvious that using color-based features does not seem promising for most medical image retrieval systems. Since then, some promising approaches to the retrieval problem were reported, among them [Schmid & Mohr, 1997] or [Ravela & Manmatha, 1998], but their applicability to medical image retrieval is still to be shown.

Furthermore, queries of diagnostic relevance include searching for organs, their relative locations and other distinct features like morphological appearances. Therefore, common retrieval systems cannot guarantee a meaningful query completion when used within medical context [Lehmann et al., 2000]. In contrast to this, the IRMA system - a joint project between three RWTH Aachen - University of Technology institutes and Philips Medical Systems - is being developed for use in daily clinical routine. This paper deals with the first important step within IRMA image processing, namely image classification. In the next Section, we will briefly explain the general IRMA architecture, motivating the need for image classification. After describing the radiograph database and the feature analysis in Section 3, we will give an overview of Simard's tangent distance [Simard et al., 1993], which we use in our classifier. The image distortion model, which - despite its simplicity - proved to be very effective in our experiments (especially in conjunction with tangent distance), is proposed in Section 5. After that, we describe our kernel density based classifier within which we use the described distance measures. Before drawing conclusions and giving an outlook to future work, we discuss our results in Section 7.

2 The IRMA-System: An Overview

From the medical point of view there exist three major applications for automated content based image retrieval [Lehmann et al., 2000] :

- (1) automatic retrieval of relevant images for follow-up studies within a PACS,
- (2) searching for representative images of known diseases and
- (3) scientific and educational studies on X-ray patterns.

In contrast to common approaches to image retrieval, the IRMA concept is based on a strict logical and algorithmic separation of the following steps to enable complex image content understanding:

- image-categorization (based on global features)
- image-registration (in geometry and contrast)
- feature extraction (based on local features)
- feature selection (category and query dependent)
- indexing (multiscale blob-representation)
- identification (incorporate a-priori knowledge)
- retrieval (on blob-level)

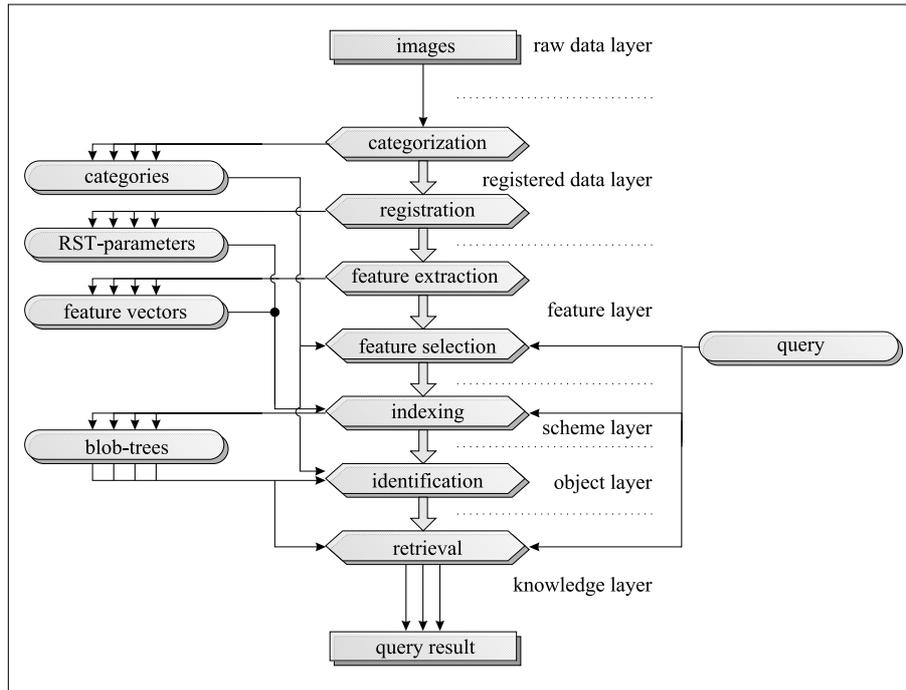


Figure 1: The IRMA architecture

To enable complex queries for medical purpose, the information retrieval system must be familiar with the class of a given image prior to query processing, as this information is of great interest for the following IRMA steps. For example, searching a pulmonal tumor in a skull radiograph is senseless (as - by definition - a pulmonal tumor is always located in the lungs), and ultrasound images need different processing than radiographs (as the characteristics of an ultrasound image greatly differ from those of a radiograph). Thus, if a radiologist is searching the image database for all radiographs showing a pulmonal tumor, the IRMA system only processes radiographs which are classified as 'chest' (or have a posterior probability for 'chest' that is higher than a user-defined treshold). On all pictures fulfilling this constraints, the (probably computational more expensive) search for tumors is done, for instance by using local textural features as proposed in [Vogelsang et al., 1997] or statistical classifiers such as proposed in [Dahmen et al., 2000]. The categorization step therefore not only reduces the computational complexity needed to answer an IRMA query, it will also most probably reduce the 'false-alarm'-rate of the system, improving its precision.

According to the deliverable of the USINT working group of the EurIPACS AIM project [Wendler et al., 1994], we define three major classes: image modality (physical), anatomic region (anatomical) and image orientation (technical). In a first step, we distinguish six anatomic regions: (1) abdomen, (2) limbs, (3) breast, (4) skull, (5) chest and (6) spine. These instances build subclasses resulting in hierarchically structured IRMA-categories. While modern DICOM

imaging devices provide information required for image classification, automatic content based classification is required for fast archiving of images acquired by film-based modalities such as radiographs. Once the class of a given image has been determined using global features, subsequent IRMA processing steps can use this information to extract problem specific features needed to answer complex queries. As classification is not necessarily unique (a chest radiograph might be labelled ‘chest’ and ‘spine’ at the same time), we call this step ‘categorization’ within the IRMA system. Thus, each image can be linked to several categories and the likelihood for each of these is also stored in the IRMA database. Therefore, classifiers used for categorization should be rather sensitive than specific.

Before discussing image categorization, we will give a brief overview of the IRMA system (see Figure 1, detailed information can be found in [Lehmann et al., 2000]): After categorization, the image is registered to a prototype which has been previously defined by an expert or by a statistical data analysis [Dahmen et al., 2000, Dahmen et al., 1999]. In the following feature extraction step we distinguish between so called ‘category-free’ features (which are suitable for all categories, i.e. a gradient image) and ‘category-specific’ features, (i.e. segmentation of the ribs in a chest radiograph [Vogelsang et al., 1998]). In the feature selection step, appropriate features for a given query are chosen. One possibility to do this is performing a linear discriminant analysis [Duda & Hart, 1973, pp. 114-123], which proved to be very efficient in first experiments [Dahmen et al., 1999]. In the indexing step, a compact representation of the given query image and the features extracted is created. Based on each set of feature images, the query image is segmented into relevant regions. Region representation (at multiple scales) will then be done via blobs [Carson et al., 1997]. This hierarchical multiscale approach will allow the user to retrieve from entire images as well as from regions of interest. The blob-identification step might be useful for queries concerning details defined within organs or other objects in an image. In the final retrieval step, the query is processed via suitable distance measures defined on the entire image or on blob-level respectively.

3 Image Database & Feature Analysis

The IRMA database we use for our experiments consists of 110 abdomen, 706 limbs, 103 breast, 110 skull, 410 chest and 178 spine radiographs, summing up to a total of 1617 images. The data is secondary digital, that is it has been scanned from conventional film-based radiographs. All images were scanned using 256 gray levels, with the image sizes ranging from about 200×200 pixels (e.g. a radiograph of a single finger) to 2000×2000 pixels (e.g. a chest radiograph). The anonymized images reflect the distribution of images in the Department of Diagnostic Radiology and were labelled by an expert (see Figure 2). Although each image is labelled with an 8-digit IRMA category code, we only concentrate on the six anatomic regions as defined in Section 2. Nevertheless, radiograph classification is a hard problem:



Figure 2: Example radiographs taken from the database, scaled to a common, square size. Top-left to bottom-right: abdomen, limbs, breast, skull, chest and spine.

On the one hand, the qualities of radiographs vary considerably and there is a great within-category variance (as caused by different doses of X-rays, varying orientations, images with and without pathologies, changing scribor position etc.). On the other hand, there is a strong visual similarity between many images of the classes abdomen and spine (cp. Figure 2). In our experiments we make use of *appearance based pattern recognition*, that is we interpret each pixel of an image as a feature. This way, all the information contained in an image is used for classification. The only preprocessing we do is downscaling the radiographs to 32×32 pixels (note that all images (and thus the according feature vectors) must be of the same size to allow for a meaningful image comparison). Our experiments showed, that this step speeds up the system significantly without notably increasing the classification error rate (cp. Section 7).

Although invariances play an important role for classifying radiographs, we do not extract invariant features. Instead, we incorporate these invariances in the classification algorithm itself. This is done by using distance measures that are - for instance - invariant to transformations like image rotation, axis deformations, scaling or varying image brightness. In the following, we will denote images of dimension $I \times J$ by

$$x = \{x_{ij} \in \mathbb{R}^+\}, i = 1, \dots, I \text{ and } j = 1, \dots, J \quad (1)$$

with x_{ij} being the grey level at pixel (i, j) and we will present invariant distance measures which will replace the (squared) Euclidean distance

$$D(x, \mu) = \|x - \mu\|^2 \quad (2)$$

between the given image x and a reference image μ .

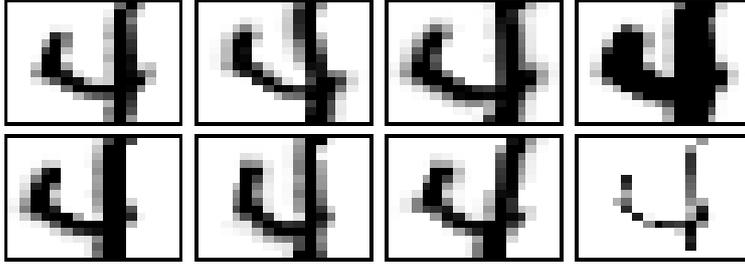


Figure 3: Example images generated via tangent approximation. Original image is at top-left.

4 Simard's Tangent Distance

In 1993, SIMARD et al. proposed an invariant distance measure called ‘tangent distance’, which proved to be especially effective for optical character recognition. The authors pointed out that reasonably small transformations of certain objects (like characters or, as in this work, radiographs) do not change their class. Simple distance measures like the Euclidean distance do not account for this, instead they are very sensitive to transformations like scaling, translation, rotation or axis deformations. When a pattern x is transformed (e.g. scaled) with a transformation $t(x, \alpha)$ which depends on a parameter α (e.g. the scaling factor), the set of all transformed patterns

$$M_x = \{t(x, \alpha) : \alpha \in \mathbb{R}\} \subset \mathbb{R}^{I \times J} \quad (3)$$

spans a one-dimensional curve in the input space. If there are L such transformations with parameters α_l , M_x is a manifold of at most L dimensions. The distance between two patterns can now be defined as the minimum distance between their according manifolds, being truly invariant with respect to the transformations t_1, \dots, t_L . Unfortunately, computation of this distance is a hard optimization problem and the manifolds needed have no analytic expression in general. Therefore, small transformations of a pattern x can be approximated by a tangent subspace \hat{M}_x to the manifold M_x at the point x . Those transformations can be obtained by adding to x a linear combination of the vectors $T_l(x)$ that span the tangent subspace. Thus, the manifold M_x can be first-order approximated by:

$$\hat{M}_x = \left\{ x + \sum_{l=1}^L \alpha_l \cdot T_l(x) : \bar{\alpha} \in \mathbb{R}^L \right\} \subset \mathbb{R}^{I \times J} \quad (4)$$

with $\bar{\alpha}$ being the vector-notation of the α_l . Now, the single-sided tangent distance $D_T(x, \mu)$ is defined as

$$D_T(x, \mu) = \min_{\bar{\alpha}} \left\{ \left\| x + \sum_{l=1}^L \alpha_l \cdot T_l(x) - \mu \right\|^2 \right\} \quad (5)$$

The so-called tangent vectors $T_l(x)$ can be computed using simple finite differences between the original image x and a reasonably small transformation of x [Simard et al., 1993]. Example images that were computed using (4) are shown in Figure 3 (we chose digits there, as they are

especially suited to demonstrate the effects of tangent approximation). One can also define a double-sided tangent distance, but this dramatically increases the computational complexity without yielding a significant improvement [Simard et al., 1993]. In our experiments, we computed the tangent vectors for translations (2), rotation, scaling and axis deformations (2) as proposed by Simard, but replaced the tangent vector for ‘line thickness’ by a ‘brightness’ tangent vector. All elements of this vector were set to a constant value to model varying image brightness, as caused by different doses of X-rays. These seven tangent vectors were used in a kernel density based classifier [Duda & Hart, 1973, pp. 61-62] for radiograph classification. Conceptionally, the single sided tangent distance is computed as follows:

- 1) compute tangent vectors for the observation x
- 2) compute an orthonormal basis for the tangent subspace (using a singular value decomposition [Press et al., 1992, pp. 59-67]).
- 3) compute the projection $\hat{\mu}(x)$ of a reference image μ into the tangent subspace of the observation x
- 4) compute $D_T(x, \mu) = D(\hat{\mu}(x), \mu)$

Note that steps 1) and 2) can be computed in advance, only steps 3) and 4) have to be done while classifying. Given that the tangent vectors are orthogonal, this can be done efficiently in a single step by computing

$$D_T(x, \mu) = \|x - \mu\|^2 - \sum_{l=1}^L \frac{[(x - \mu)^t \cdot T_l(x)]^2}{\|T_l(x)\|^2} \quad (6)$$

The tangent distance is a very effective means to compensate small global transformations of an image. In the following we will propose a simple, yet very effective image distortion model for local variations. We will show that both approaches work very well, but that the best results are obtained by combining both.

5 The Image Distortion Model

The last conceptional step in the computation of the tangent distance still requires the calculation of the (squared) Euclidean distance between an image μ and its projection $\hat{\mu}(x)$ into the tangent subspace of a second image x . Although small global transformations have been compensated by the projection step, this distance is still highly sensitive to *local* transformations of the images, e.g. caused by noise (which is typical for radiographs). We therefore propose the following image distortion model:

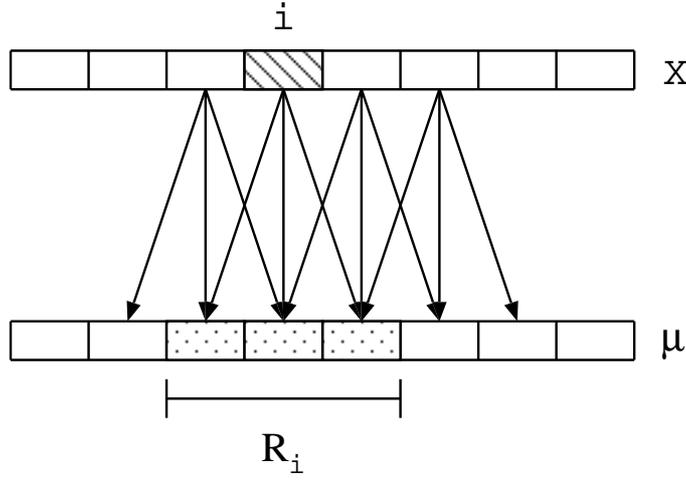


Figure 4: One-dimensional example of the distortion model with $r=1$.

When calculating the distance between two images x and μ we allow for local deformations, i.e. we do not compute the squared error between a pixel (i, j) in x and its counterpart in μ , but look for the ‘best-fitting’ pixel in μ within a certain neighbourhood R_{ij} (see Figure 4):

$$D_{dist}(x, \mu) = \sum_{i=1}^I \sum_{j=1}^J \min_{(i', j') \in R_{ij}} \|x_{ij} - \mu_{i'j'}\|^2 \quad (7)$$

for images with dimension $I \times J$. We typically chose R_{ij} to be square. Obviously, this unrestricted distortion approach can model wanted as well as unwanted (i.e. meaningless) transformations. Nevertheless, an appropriate choice of R_{ij} leads to a significant improvement of radiograph classification.

5.1 An extended Distortion Model

From Eq. (7) it becomes evident that with increasing neighbourhood R_{ij} the transformations realized by the distortion model violate the assumption that the class-membership of the original input image equals that of the transformed input image. In fact, the distortion distance between almost any two images can be reduced to a value near zero by increasing R_{ij} , leading to a significant increase in classification error. To compensate for this, we introduce a cost function $C(i, i', j, j')$, which models the costs for deforming a pixel x_{ij} in the input image to a pixel $\mu_{i'j'}$ in the reference image:

$$D_{dist}(x, \mu) = \sum_{i=1}^I \sum_{j=1}^J \min_{(i', j') \in R_{ij}} \{ \|x_{ij} - \mu_{i'j'}\|^2 + C(i, i', j, j') \} \quad (8)$$

We propose two methods to estimate $C(i, i', j, j')$:

- choose $C(i, i', j, j')$ empirically, e.g. by using a weighted Euclidean distance between pixels (i, j) and (i', j') . This way, small local transformations are preferred to (most probably unwanted) long-range pixel transformations.
- learn $C(i, i', j, j')$ by using training samples and a maximum-likelihood approach. That is, do meaningful transformations in training and choose $C(i, i', j, j')$ using relative frequencies of possible transformations. The more often a transformation was performed in training, the lower its cost.

6 Classification

For classification, we use the distance measures described above in a kernel density based classifier (KD) and compare the obtained results with that of a simple 1-nearest neighbour classifier (1-NN). As there are only 1617 images available so far, splitting the data into a training and a testing set is not advisable. To obtain meaningful results, we therefore make use of a *leaving-one-out* approach. Given an image x to classify, we use all the remaining images as reference patterns and compute the class conditional probability $p(x|c)$ for the image x given class c as follows:

$$p(x|c) = \frac{1}{N_c} \sum_{n=1}^{N_c} \frac{1}{A_{c,\gamma}} \exp\left(-\frac{d(x, x_{nc})}{\sigma_c^2 \cdot \gamma}\right) \quad (9)$$

with N_c being the number of reference images of class c , x_{nc} being the n -th reference pattern of class c , class specific standard deviation σ_c^2 and $d(x, x_{nc})$ being one of the proposed distance measures. To compensate for the fact that variances are usually underestimated if only few training samples are available, we multiply the estimated variances with a factor $\gamma > 1$. Strictly speaking, the normalization factor $A_{c,\gamma}$ depends on the class c , however, the dependency is weak and therefore neglected in our experiments. Classification is then done using the Bayesian decision rule [Duda & Hart, 1973, pp. 10-39], i.e. by computing:

$$x \mapsto r(x) = \underset{c}{\operatorname{argmax}} \{p(c) p(x|c)\} \quad (10)$$

Modelling the prior probabilities $p(c)$ by using relative frequencies and neglecting the normalization factor, we obtain

$$x \mapsto r(x) = \underset{c}{\operatorname{argmax}} \left\{ \frac{N_c}{N} \frac{1}{N_c} \sum_{n=1}^{N_c} \exp\left(-\frac{d(x, x_{nc})}{\sigma_c^2 \cdot \gamma}\right) \right\} \quad (11)$$

Table 1: Leaving-one-out error rates for varying distance measures without tresholding

Distance Measure	Error Rate [%]	
	1-Nearest Neighbour	Kernel Densities
Euclidean Distance	18.1	16.4
Tangent Distance	15.3	14.8
Distortion Model	16.5	14.7
Distorted Tangent Distance	14.2	12.5

where N is the total number of reference images. As the term $(1/N)$ does not depend on c , it can be neglected in the maximization step. Thus, we finally obtain:

$$x \mapsto r(x) = \operatorname{argmax}_c \left\{ \sum_{n=1}^{N_c} \exp \left(-\frac{d(x, x_{nc})}{\sigma_c^2 \cdot \gamma} \right) \right\} \quad (12)$$

In a last step, we compute posterior probabilities $p(c|x)$ for each class c given the image x by computing:

$$p(c|x) = \frac{p(c) \cdot p(x|c)}{\sum_{c'=1}^C p(c') \cdot p(x|c')} \quad (13)$$

with C being the number of classes. We can now sort the possible image classes for x by decreasing probability and obtain a class ranking list. In the experiments, besides measuring the classification error rate, we also compute the ‘ranking’ of an image. If the correct class has rank 1, the image was classified correctly. If this is not the case, we still distinguish between a rank 2 error and - for instance - a rank 6 error (being the worst case). To reduce the computational complexity of the following IRMA processing steps, the mean ranking over all images should be close to 1, as in this case the number of hypotheses expanded by the IRMA system can be reduced.

7 Results

For our experiments, we scaled down the radiographs to a standard size of 32×32 pixels. This can be done without a significant change in classification error rate, but leads to a considerable system speedup. Computing a simple 1-NN (performing the *leaving-one-out* approach as described in Section 6) on the radiographs with a size of 320×320 pixels gives a classification error of 18.0%, requiring about 30 CPU seconds on a 500MHz Digital ALPHA CPU to classify a single image. Downscaling the images to the proposed size of 32×32 pixels, we obtained an error rate of 18.1% (with an average image ranking of 1.30), requiring about 0.4 CPU seconds.

Table 2: Leaving-one-out error rates for varying distance measures with tresholding ($S=3500$)

Distance Measure	Error Rate [%]	
	1-Nearest Neighbour	Kernel Densities
Euclidean Distance	14.8	14.2
Tangent Distance	13.7	12.9
Distortion Model	14.5	13.2
Distorted Tangent Distance	10.6	10.3

Having chosen the image size, we used the single-sided tangent distance for radiograph classification. As can be seen in Table 1, this reduces the KD error rate from 16.4% to 14.8%. We then started experiments with the image distortion model, using $C(i, i', j, j') = 0$. Surprisingly, with an error rate of 14.7% the result of this simple distortion model is even slightly better than that obtained by using tangent distance. In another experiment we tried to find out whether the gains of both approaches are additive. Indeed, combining both distance measures (i.e. computing $D_{dist}(\hat{\mu}(x), \mu)$ in step 4) of tangent distance instead of $D(\hat{\mu}(x), \mu)$) reduced the error rate from 14.8% to 12.5% ('distorted tangent distance'). Figure 5 shows the achieved error rates with respect to the size of a square neighbourhood R of dimension $(2r + 1) \times (2r + 1)$. The best result of 12.5% was obtained using $r = 0.7$ (using linear interpolation between pixels).

In another experiment we restricted the maximum local distance between two image pixels by a threshold S . Note that the maximum contribution of a pixel to any of the proposed distance measures is $255 \cdot 255 = 65.025$, as the radiographs are 256-grayscale images. Thus, a single pixel may have a significant contribution to the total distance, so that a few distorted pixels (as caused by noise or changing scribor position) can lead to a misclassification. By restricting this contribution to a maximum value S ('local tresholding') we can compensate for this effect. Again, this very simple method yields a significant improvement in classification error rate, namely an error rate of 10.3% (see Table 2), with an average image ranking of 1.18.

Analysing the remaining errors we found out that many misclassifications could be easily avoided by taking into consideration the original image aspect ratios (by downscaling the images to a standard size we lose this information). To compensate for this we introduced an aspect ratio penalty term, based on the difference in aspect ratio between the given image and the reference image. This penalty term reduced the classification error from 10.3% to 8.9%. We then chose $C(i, i', j, j')$ to be a weighted Euclidean distance between pixels (see Section 5.1), further reducing the error rate from 8.9% to 8.6% with ranking 1.17 (with the class-specific error rates ranging from 27.3% for abdomen to 3.4% for chest). Note that the error rate drops to less than 2% if we consider the top three classes, i.e. the three classes with highest probabilities. Thus, despite the error rate of 8.6%, the proposed algorithms are already suitable for use within the IRMA-system.

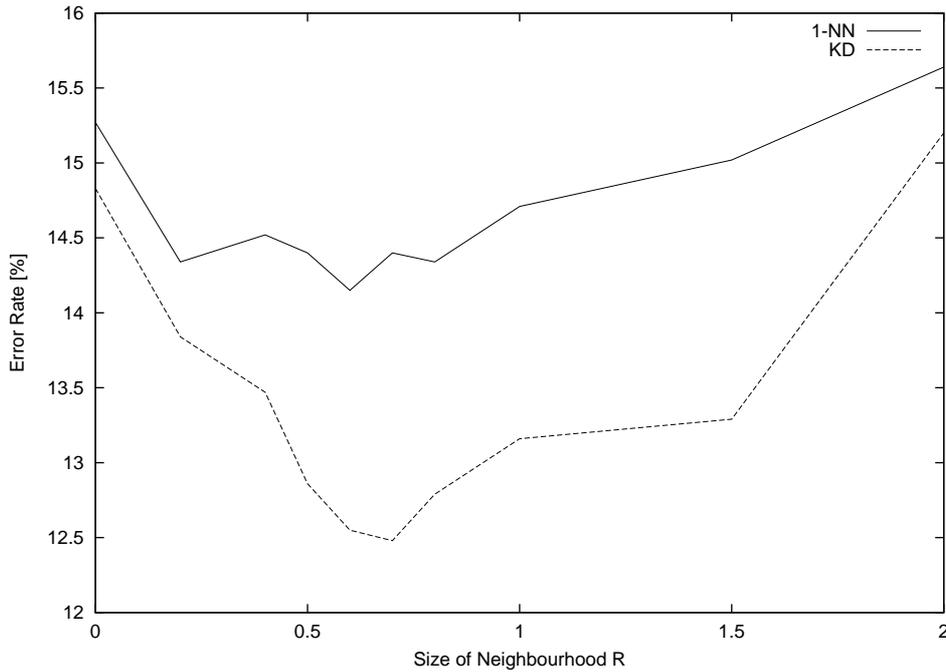


Figure 5: Error rates for distorted tangent distance with respect to size of neighbourhood R , without local thresholding

In a final experiment, we analysed the different distance measures we discussed above with respect to their invariance properties, given a transformation t . In our experiments, we chose t to be a translation and calculated the distance between a shifted version of a radiograph and the original image as well as the distance to radiographs from competing classes. As we can see in Figure 6, the Euclidean distance is highly sensitive to image translations. On the other hand, the tangent distance (Figure 7) can nearly compensate 1 pixel shifts and yields small distances up to 2-3 pixels shifts. As expected, the distortion model with $r = 1$ (as shown in Figure 8) can fully compensate 1 pixel shifts, yet with R increasing, the distances to competing classes get smaller rapidly (see Figures 8 and 9). Thus, large neighbourhoods may lead to bad classification results.

7.1 On Cooccurrence Matrices

The use of cooccurrence matrices [Haralick et al., 1973] is often considered to be helpful for content based medical image retrieval. However, our experiments on radiograph classification do not support this thesis. In two experiments, we used global cooccurrence matrices for feature analysis within a synergetic classifier [Haken, 1991] and within a kernel density based classifier. In both cases, we were not able to obtain classification error rates below 29%. Apparently, cooccurrence matrices do not provide discriminative features for radiograph classification. Nevertheless they might still be useful for the following IRMA processing steps, e.g. to detect tumors within a (previously categorized) radiograph. In this case, cooccurrence matrices would be computed from small parts of the image, not from the complete image [Vogelsang et al., 1997].

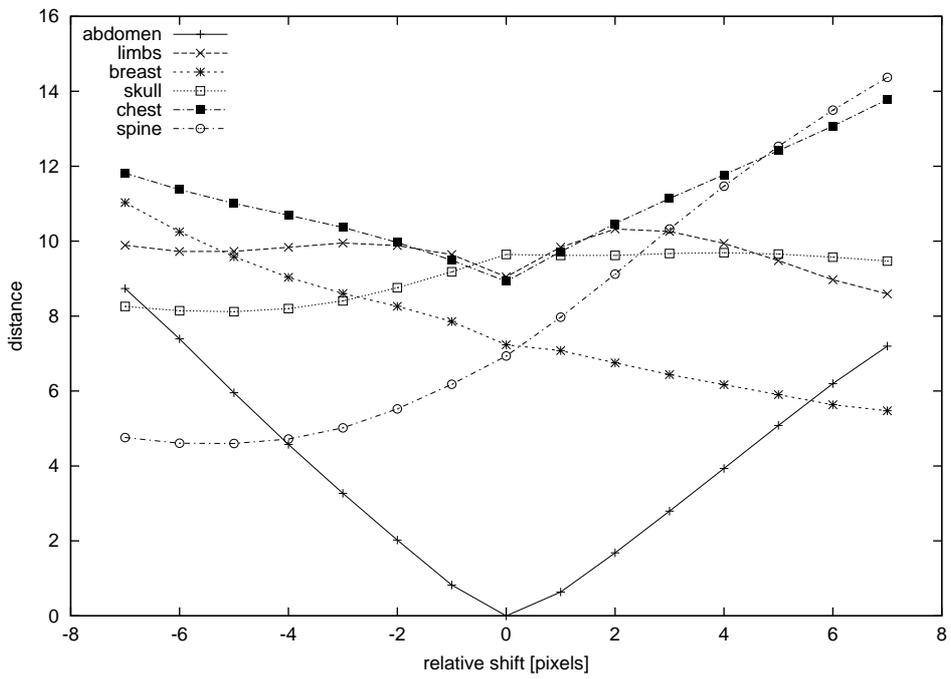


Figure 6: Behaviour of Euclidean distance with respect to image shifts

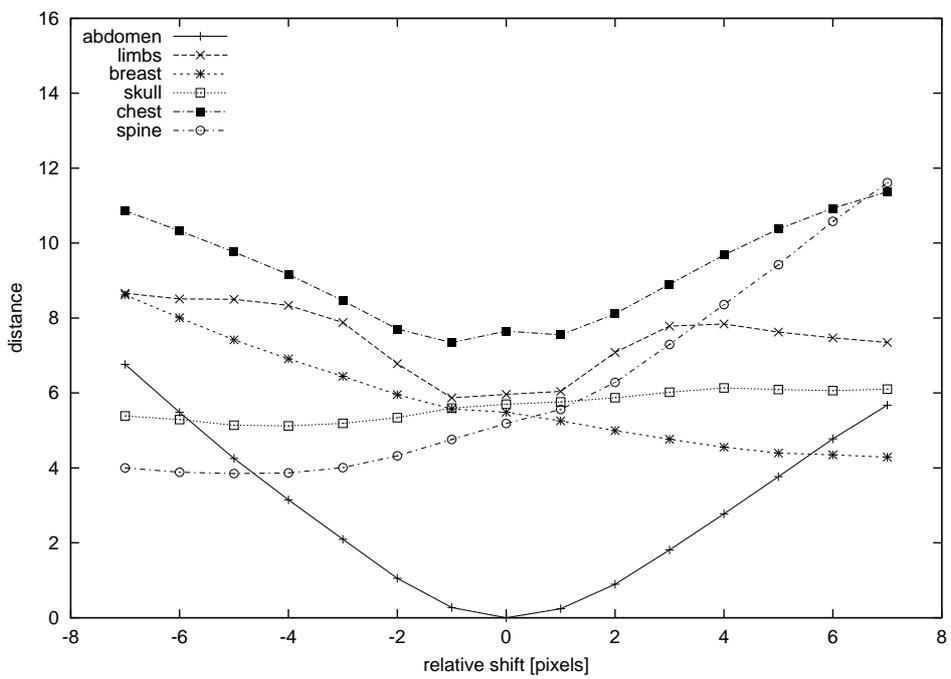


Figure 7: Behaviour of tangent distance with respect to image shifts

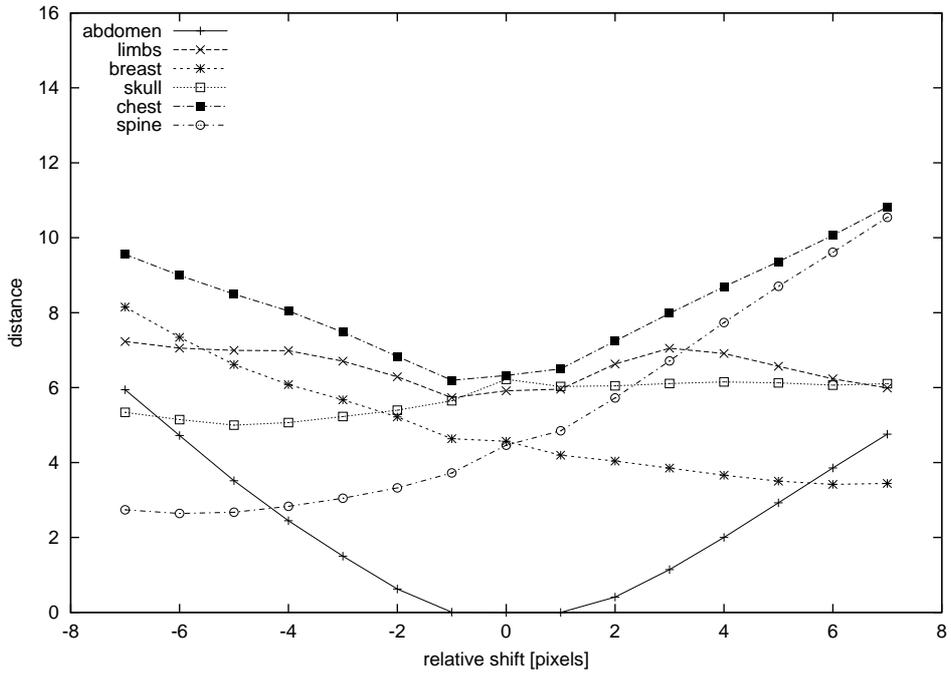


Figure 8: Behaviour of distortion distance with respect to image shifts, using a neighbourhood with $r = 1$.

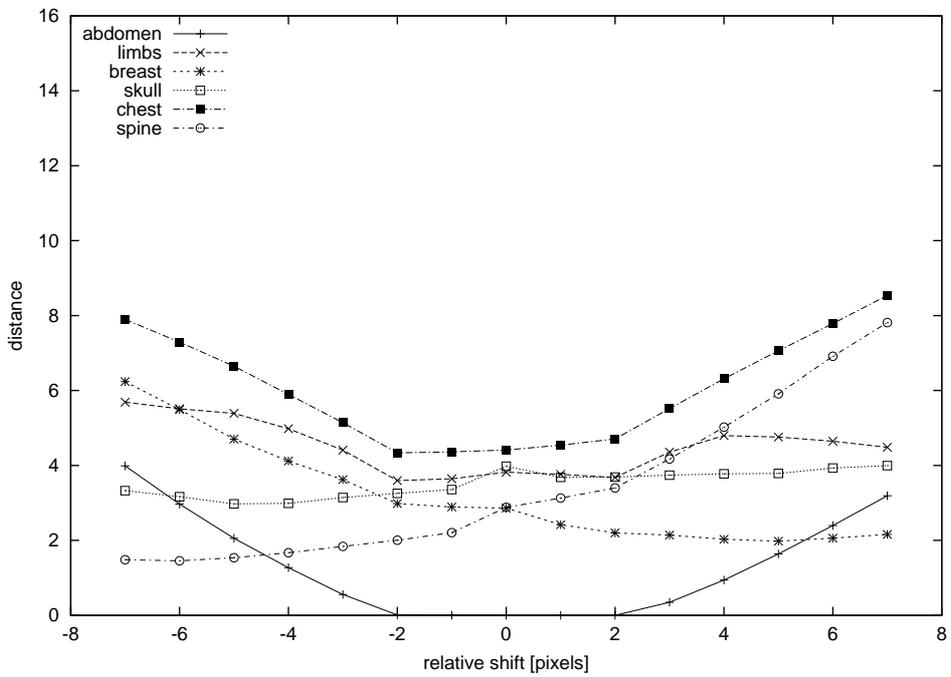


Figure 9: Behaviour of distortion distance with respect to image shifts, using a neighbourhood with $r = 2$.

8 Conclusions

In this paper we presented our approach to classifying radiographs within the IRMA system. We proposed a new image distortion model and proved its efficiency by applying it to 1617 radiographs coming from daily routine. Although the distortion model alone is very effective, we showed that it can be used to improve Simard's tangent distance in this application. The best classification error of 8.6% was achieved by computing this 'distorted tangent distance', which means a relative improvement of 52% with respect to the Euclidean distance nearest neighbour error rate of 18.1%. We are currently working on the improvement of our image distortion model. This can be done by a better choice of the cost function $C(i, i', j, j')$, e.g. by learning it from the training data. We will furthermore use additional constraints to restrict possible image distortions to 'meaningful' ones.

References

- [Carson et al., 1997] C. Carson, S. Belongie, H. Greenspan, J. Malik, "Color- and Texture-Based Image Segmentation using EM and its Application to Image Querying and Classification", to appear in: IEEE Pattern Analysis and Machine Intelligence.
- [Dahmen et al., 1999] J. Dahmen, R. Schlüter, H. Ney, "Discriminative Training of Gaussian Mixtures for Image Object Recognition", in W. Förstner, J. Buhmann, A. Faber, P. Faber (eds.): Proceedings of the 21. Symposium German Association for Pattern Recognition, pp. 205-212, Bonn, September 1999.
- [Dahmen et al., 2000] J. Dahmen, K. Beulen, M. Gueld, H. Ney, "A Mixture Density Based Approach to Object Recognition for Image Retrieval", Proceedings of the 6th International RIAO Conference on Content-Based Multimedia Information Access, Paris, France, April 2000, this issue.
- [Duda & Hart, 1973] R. Duda, P. Hart, *Pattern Classification and Scene Analysis*, John Wiley & Sons, New York, 1973.
- [Faloutsos et al., 1994] C. Faloutsos, R. Barber, M. Flickner, J. Hafner et al., "Efficient and Effective Querying by Image Content", Journal of Intelligent Information Systems, Vol.3, pp. 231-262, 1994.
- [Haken, 1991] H. Haken, *Synergetic Computers and Cognition*, Springer Verlag, New York, pp. 36-50, 1991.
- [Haralick et al., 1973] R. Haralick, K. Shanmugam, I. Deinstein, "Textural features for Image Classification", IEEE Transactions on Systems, Man and Cybernetics, Vol.3, No.6, pp. 610-621, 1973.

- [Lehmann et al., 2000] T. Lehmann, B. Wein, J. Dahmen, J. Bredno, F. Vogelsang, M. Kohnen, "Content-based Image Retrieval in Medical Applications: A Novel Multi-step Approach", Proceedings of the International Society for Optical Engineering (SPIE), Vol. 3972(32), 2000, in press.
- [Lou et al., 1997] S. Lou, D. Hoogstrate, H. Huang, "An Automated PACS Image Acquisition and Recovery Scheme for Image Integrity Based on the DICOM Standard", Computerized Medical Imaging and Graphics, Vol.21, No.4, pp. 209-218. 1997.
- [Press et al., 1992] W. Press, S. Teukolsky, W. Vetterling, B. Flannery, *Numerical Recipes in C*, University Press, Cambridge, 1992.
- [Ravela & Manmatha, 1998] S. Ravela, R. Manmatha, "On Computing Global Similarity in Images", Proceedings of the IEEE Workshop on Applications of Computer Vision (WACV), pp. 82-87, Princeton, NJ, October 1998.
- [Schmid & Mohr, 1997] C. Schmid, R. Mohr, "Local Grayvalue Invariants for Image Retrieval", IEEE Transactions on Pattern Recognition and Machine Intelligence, Vol. 19, No. 5, pp. 530-535, May 1997.
- [Simard et al., 1993] P. Simard, Y. Le Cun, J. Denker, "Efficient Pattern Recognition Using a New Transformation Distance," S.J. Hanson, J.D. Cowan, C.L. Giles (eds.): *Advances in Neural Information Processing Systems 5*, Morgan Kaufmann, San Mateo CA, pp. 50-58, 1993.
- [Vogelsang et al., 1997] F. Vogelsang, F. Weiler, B. Wein, M. Kilbinger, R. Günther, "Image Content Analysis using Textural Information and Synergetic Classifiers", Proceedings of the European Congress of Radiology ECR 97, pp. S 291, Vienna, Austria, 1997.
- [Vogelsang et al., 1998] F. Vogelsang, F. Weiler, J. Dahmen, M. Kilbinger, B. Wein, R. Günther, "Detection and Compensation of Rib Structures in Chest Radiographs for Diagnose Assistance", Proceedings of the International Symposium on Medical Imaging 1998 (SPIE), pp. 774-785, Vol. 3338-1, San Diego, 1998.
- [Wendler et al., 1994] T. Wendler, B. Wein, R. van den Brock, Deliverable USINT in the EurIPACS-Project A1009, Technical Report, AIM, Ref.XIII, European Commission, Brussels, February 1994.