

Structural Comparison of Mammograms

Byung-Woo Hong
Medical Vision Lab
University of Oxford
Oxford, UK

Michael Brady
Medical Vision Lab
University of Oxford
Oxford, UK

Abstract

This paper presents a robust algorithm for the comparison of mammogram pairs. Salient regions are extracted in a topographic way. An integral invariant representation of shape, in combination with area and distance measures, are used for establishing their correspondences. The experimental results demonstrate that the algorithm can provide a useful tool for a Computer-Aided Diagnosis system in mammography.

1 Introduction

Breast cancer is the leading cause of death from cancer among women in the West and is responsible for one in five of all female cancers. In the West, one in nine women will develop breast cancer at some time in their life time. Approximately 13,000 women die from breast cancer in England and Wales alone each year. The earlier the diagnosis, the better the prognosis, primarily because it lessens the risk of metastasis to a vital organ. The UK breast screening programme (BSP) was the first of its kind in the world and has played a major role in reducing mortality from breast cancer. All screening programmes are based on x-ray mammography. Unfortunately, approximately 24% of cancers are missed at screening, primarily because of: the complexity and variability of breast anatomy throughout the population, the subtlety of clinically significant signs, variations in image contrast resulting from imaging parameters (e.g. exposure time) [4], and because of the huge workload falling on the BSP (1.5 million women per year have mammograms). A mammogram may be inspected for just a few seconds. The workload is increasing sharply as a result of the decision to require two different views of each breast, and because there is pressure to reduce the interval between screening rounds, since the number of “interval cancers” seen at a screening round but not seen/missed at the previous screening round are far higher than originally predicted.

There is increasing interest in computer-aided detection (CAD) to support radiologists in examining mammograms. To date, CAD systems have enjoyed considerable success in detecting microcalcifications; but there do not yet exist CAD systems that can detect malignant masses and other abnormalities with sufficient sensitivity and specificity. A fundamental observation, the starting point for this paper, is that even highly-skilled radiologists often cannot diagnose an abnormality on the basis of a single mammogram. Indeed, it is accepted that there is no absolute basis for the detection of abnormalities. Instead, breast radiologists are trained to **compare** a mammogram with others. In **temporal** comparisons, radiologists compare a current mammogram with the corresponding

one from the previous round and look for signs of significant change. In **bilateral** comparisons, radiologists compare the mammograms of the left and right breasts and look for “asymmetries” and “architectural distortions” between the images, though both of these terms remain frustratingly vague. Temporal comparison has the advantage that nominally the same breast is imaged: nominally, because of the changes wrought by involution at the menopause, HRT, and because an interval cancer may be visible in the later image but not the first. We propose here an algorithm for temporal comparison of mammograms. More details and examples may be found in [1]. It should be understood that even if there is no significant change, the two images may appear quite different, even if (though it is unlikely) the image acquisition parameters are the same. The reason is that much of the difficulty of mammogram comparison stems from the tight compression of the breast (typically 150N) while it is imaged. A slight difference in compression can cause connective stroma to appear/disappear, changing the image markedly. Moreover, since tumours are often attached by tendon-like spiculations to the inside surface of the breast, and because tumours are biomechanically dense, a small change in compression can cause a substantial displacement of a tumour region.

Differences in breast compression (a 1cm change is typical) mean that, in practice, rigid and affine transformations between the images are insufficient. It has been observed [6] that the breast boundary provides a useful basis for aligning the images, in their case using a thin-plate spline. However, it has also been demonstrated that this is not sufficient - internal “landmarks” are necessary to supplement the boundary. Some authors have proposed using the pectoral muscle, which is clearly visible on mediolateral oblique mammograms; but the relative lack of control of arm position during image formation means that the alignment is often worse if the pectoral muscle is used to further constrain alignment. If we are to use internal “landmarks”, what, in view of the immense variability in anatomy, should they be? Some authors have suggested locally bright points such as calcifications; but these are (happily) only present in a minority of mammograms. That is, a set of points defined by local filtering (e.g. “corner” points) do not make anatomical sense. However, **every** mammogram contains regions of dense tissue, which correspond either to normal parenchymal tissue, to lesions (benign or malignant), and, because they are projected x-ray images, overlaps of such regions with additional brightness. These considerations inform the algorithm that we have developed: (a) extract from each of a pair of mammogram a manageable set of “salient” regions; (b) establish a set of correspondences between such regions using their outline shapes; and then (c) determine if there are unpaired regions or regions whose shape/size has changed significantly.

2 Structural Comparison

The problem of comparing mammograms is cast in the framework of tree matching. Mammographic features are characterised by a particular subset of iso-contours. Their enclosure relationships are represented by the inclusion tree that is constructed from the extracted iso-contours. A tree structure is a useful representation to identify “salient” regions and it can represent the hierarchical decompositions of mammographic features effectively. The algorithms for constructing the inclusion tree based on iso-contours and detecting salient regions in mammograms utilising the inclusion tree are presented in [2].

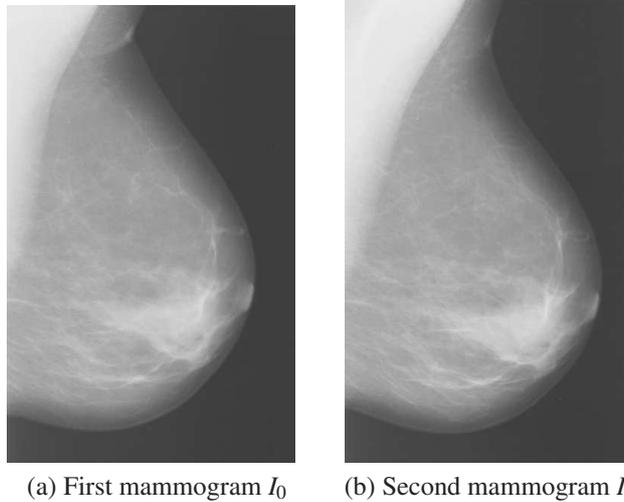


Figure 1: A pair of temporal mammograms. (a) A mammogram. (b) A mammogram of the same breast as (a), but taken at a later time.

Mammogram comparison is achieved by matching their inclusion trees, in which each node represents a contour. A match between two tree structures is generally designed to find correspondences between all of their nodes. However, in the case of mammogram comparison, it is unreasonable to employ such a *dense* correspondence since they may present different image contents. In addition to dissimilarities between corresponding mammograms, insignificant mammographic features are inevitably presented in mammograms and depict complex inner structures of the breast. It is likely that such insignificant features in one mammogram do not have any corresponding feature in the other. Consequently, it is not desirable to find correspondences for all the objects presented in a pair of mammograms. Therefore, the nodes of the contours representing the boundaries of salient regions are considered to be matched when the inclusion trees of the two mammograms are matched. The salient regions found for the pair of mammograms shown in Figure 1 are shown in Figure 2. [1] presents similar results for numerous pairs of mammograms.

2.1 Implicit Tree Matching

A correspondence between nodes in a pair of inclusion trees is obtained by comparing the contours on the nodes. We propose a cost function that measures the similarity between contours based both on their topological and geometrical properties. The shape is the primary information for measuring the similarity between contours. As a shape descriptor, an integral invariant is employed and the shape distance between contours is measured based on the similarity of their integral invariants [5]. Given two images $I_0(\mathbf{x})$ and $I_1(\mathbf{x})$, let $\Omega_0 \subset \mathbb{R}^2$ and $\Omega_1 \subset \mathbb{R}^2$ be the respective domains of the images. We aim to find a transformation $\phi : \Omega_0 \rightarrow \Omega_1$ that aligns the image $I_0(\mathbf{x}), \mathbf{x} \in \Omega_0$ to the image $I_1(\phi(\mathbf{x})), \mathbf{x} \in \Omega_0$ subject to a suitable similarity measure. This is equivalent to finding a

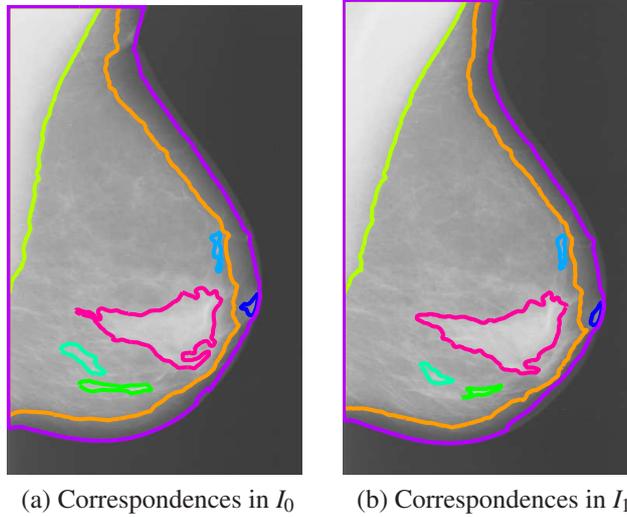


Figure 2: The correspondences between regions of interest. Each contour in (a) corresponds to the contour of the same colour in (b). (a) Contours which have corresponding contours in C_1^l among the contours in C_0^k . (b) Contours which have corresponding contours in C_0^k among the contours in C_1^l .

transformation that minimises a dissimilarity between $I_0(\mathbf{x}), \mathbf{x} \in \Omega_0$ and $I_1(\phi(\mathbf{x})), \mathbf{x} \in \Omega_0$. In the tree matching framework, the optimal transformation ϕ is determined based on correspondences between pairs of corresponding regions $R_0^i \subset I_0$ and $R_1^j \subset I_1$ where the regions R_0^i and R_1^j are precomputed as salient regions.

The iso-contour maps $CM(I_0)$ and $CM(I_1)$ are initially constructed from de-noised versions of I_0 and I_1 using anisotropic diffusion. Let C_0^i and C_1^j be iso-contours extracted at the same number of equally spaced intensities from I_0 and I_1 respectively so that $CM(I_0)$ and $CM(I_1)$ have the same resolution. The inclusion tree is then built based on the topological relationship between contours in each contour map. Distinctive image features are well characterised in the contour map and significant features are identified by analysing the inclusion tree. The contours representing the boundaries of salient regions in I_0 and I_1 are selected from $CM(I_0)$ and $CM(I_1)$ by identifying base contours in the corresponding inclusion trees. We refer to them as ‘‘salient contours’’. We denote the sets of salient contours from $CM(I_0)$ and $CM(I_1)$ by $SC(I_0)$ and $SC(I_1)$, where $SC(I_j) = \{C_j^k\}$. In order to match images I_0 and I_1 , we need to find a contour C_1^l in $SC(I_1)$ corresponding to each contour C_0^k in $SC(I_0)$. This is performed by using the inclusion trees built from $SC(I_0)$ and $SC(I_1)$.

2.2 Correspondence of Contours

In order to establish correspondences between two sets of salient contours $SC(I_0)$ and $SC(I_1)$, we define a cost function $E(C_0, C_1)$ that measures the similarity between contours

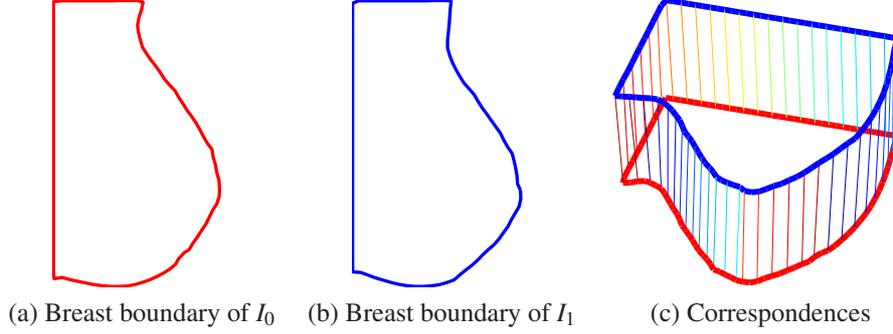


Figure 3: Correspondences between a pair of corresponding breast boundaries.

C_0 and C_1 . This cost function $E(C_0, C_1)$ is defined by:

$$E(C_0, C_1) = E_{shape}(C_0, C_1) \cdot \exp^{E_{area}(C_0, C_1) + \alpha \cdot E_{dist}(C_0, C_1)} \quad (2.1)$$

where $E_{shape}(C_0, C_1)$ is given by the symmetric shape distance $D(C_1, C_2) = E(I_1, I_2, d^*; s)$, in which the energy function E is defined by

$$E(I_1, I_2, d; s) = E_1(I_1, I_2, d; s) + E_2(d'; s) \quad (2.2)$$

$$= \int_0^1 \|I_1(s - d(s)) - I_2(s + d(s))\|^2 ds + \alpha \int_0^1 \|d'(s)\|^2 ds \quad (2.3)$$

where (I_1, I_2) are integral invariants computed for the shapes of the regions and the optimisation is performed in a dynamic programming framework [1]. The point-wise correspondences obtained by minimising the shape distance between a pair of breast boundaries are presented in Figure 3. The difference between the areas enclosed by the contours is taken into account by $E_{area}(C_0, C_1)$ which is defined by:

$$E_{area}(C_0, C_1) = 2 \cdot \frac{|A_0 - A_1|}{(A_0 + A_1)} \quad (2.4)$$

where A_0 and A_1 are the areas of the regions enclosed by contours C_0 and C_1 respectively. In the cost function, $E_{dist}(C_0, C_1)$ measures a Euclidean distance between centroids of contours C_0 and C_1 . The parameter α in Equation (2.1) is chosen according to the relative significance between the area difference and the distance measure. For a given contour C_0 in $SC(I_0)$, its corresponding contour C_1 in $SC(I_1)$ is obtained as follows:

$$C_1 = \arg \min_C E(C_0, C), \quad \forall C \in SC(I_1) \quad (2.5)$$

We introduce the notation \approx for a regional correspondence between contours. Thus, $C_0 \approx C_1$ indicates that C_0 corresponds to C_1 with respect to the cost function E .

Based on the cost function, we have developed a search algorithm for matching two sets of salient contours utilising the inclusion tree hierarchically. For each contour C_0^k in $SC(I_0)$, a correspondence $C_0^k \approx C_1^l$ is established by selecting a best matching contour

C_1^l in $SC(I_1)$ with respect to the cost function E . The matching algorithm is applied recursively to each contour C_0^k in $SC(I_0)$ by searching the inclusion tree built from $SC(I_1)$. The searching criteria for finding best correspondences between contours is constrained by their topological properties characterised by the inclusion tree. For a contour C_0^k in $SC(I_0)$, the search space for its best matching contour C_1^l in $SC(I_1)$ is restricted to the set of contours under the subtree whose root is the corresponding contour of the immediate parent of C_0^k in the inclusion tree built from $SC(I_0)$. Assuming a correspondence $C_0^k \approx C_1^l$ where $C_0^k \in SC(I_0)$ and $C_1^l \in SC(I_1)$, for a contour $C_0^{k'}$ which is an immediate descendant of C_0^k in the inclusion tree built from $SC(I_0)$, its corresponding contour is searched under the subtree whose root is C_1^l in the inclusion tree built from $SC(I_1)$. Note that a significantly strong correspondence that breaks the topological constraint is evidence for a mass. As noted earlier, soft tissue regions tend to preserve their local topology under a small change in compression. However, a tumour can move abruptly even with a slight change, possibly inducing a topological change.

Each time a correspondence is established between each contour C_0^k in $SC(I_0)$ and a contour C_1^l in $SC(I_1)$, the transformation that aligns the image I_0 to the image I_1 may be approximated. However, any correspondence that has a high matching cost is rejected. The alignment is obtained by a warping algorithm using thin-plate splines [3].

To summarise, for a given pair of temporal mammograms I_0 and I_1 of the same breast at different times, the procedures for comparing one image to the other are as follows:

1. Extract salient regions (ROI's) from the two images I_0 and I_1 .
2. Establish regional correspondences (\approx) between two sets of salient regions $SC(I_0)$ and $SC(I_1)$.
3. Establish point-wise correspondences (\sim) between each pair of corresponding regions delineated by C_0^k in $SC(I_0)$ and C_1^l in $SC(I_1)$.
4. Calculate a transformation from the image I_0 to the image I_1 by applying thin-plate splines based on the established point-wise correspondences.
5. Apply a warping to the image I_0 by a bilinear interpolation based on the obtained transformation.
6. Compare the warped image of the image I_0 and the image I_1 by subtracting one image from the other.

3 Results

Figure 1 shows a temporal pair of mammograms. Clinically, they are quite similar, though variations in size and appearance are evident. The contours representing the boundaries of the regions of interest are shown in Figure 2, colour coded to show the correspondences computed for the regions. Once regions of interest are extracted from the two images, correspondences are established based on the cost function defined in Equation (2.1): we have set $\alpha = 2$ throughout the many experiments we have done. In Figure 4 and 5, a set of corresponding pairs (first and second columns) of regions are shown, along with their

optimal point-wise correspondences (third column) where the matching cost for each pair of regional correspondence is also presented.

The transformation that is estimated based on the point-wise correspondences between the salient contours in the images I_0 and I_1 . The registration results are presented in Figure 6 where the superimposition of the original salient contours is shown in (a) in contrast to the superimposition of aligned salient contours shown in (b). The resulting deformation grid is shown in (c) and the warped image of I_0 according to the optimal deformation is presented in (d).

The algorithm has been applied to dozens of pairs of mammograms, with very encouraging results that will be analysed in a forthcoming paper.

4 Conclusions

We have developed an algorithm to facilitate the comparison of pairs of mammograms, based on matching salient contours in the two images. Salient contours provide a useful representation of significant breast structures because they effectively capture reliable features that can be used as landmarks for registration. A topographic representation using inclusion trees is the basis for efficiently and effectively finding correspondences. An integral invariant representation of shape, in combination with area and distance measures, are shown to be useful features for measuring the similarity between contours. Interpolation using thin-plate splines robustly represents the complex deformations that naturally occur in mammogram pairs. The effectiveness of the registration algorithm is currently qualitative, based on difference images. We are currently evaluating a quantitative measure of registration accuracy based on Standard Mammogram Form [4].

References

- [1] Anonymous. PhD thesis.
- [2] Anonymous.
- [3] F. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11(6):567–585, 1989.
- [4] R. Highnam and M. Brady. *Mammographic Image Analysis*. Kluwer Academic Publisher, 1999.
- [5] S. Manay, B.-W. Hong, A. Yezzi, and S. Soatto. Integral invariant signatures. *In Proc. European Conference on Computer Vision*, May, 2004.
- [6] M. Y. Sallam and K. W. Bowyer. Registration and difference analysis of corresponding mammogram images. *Medical Image Analysis*, 3(2):103–118, 1999.

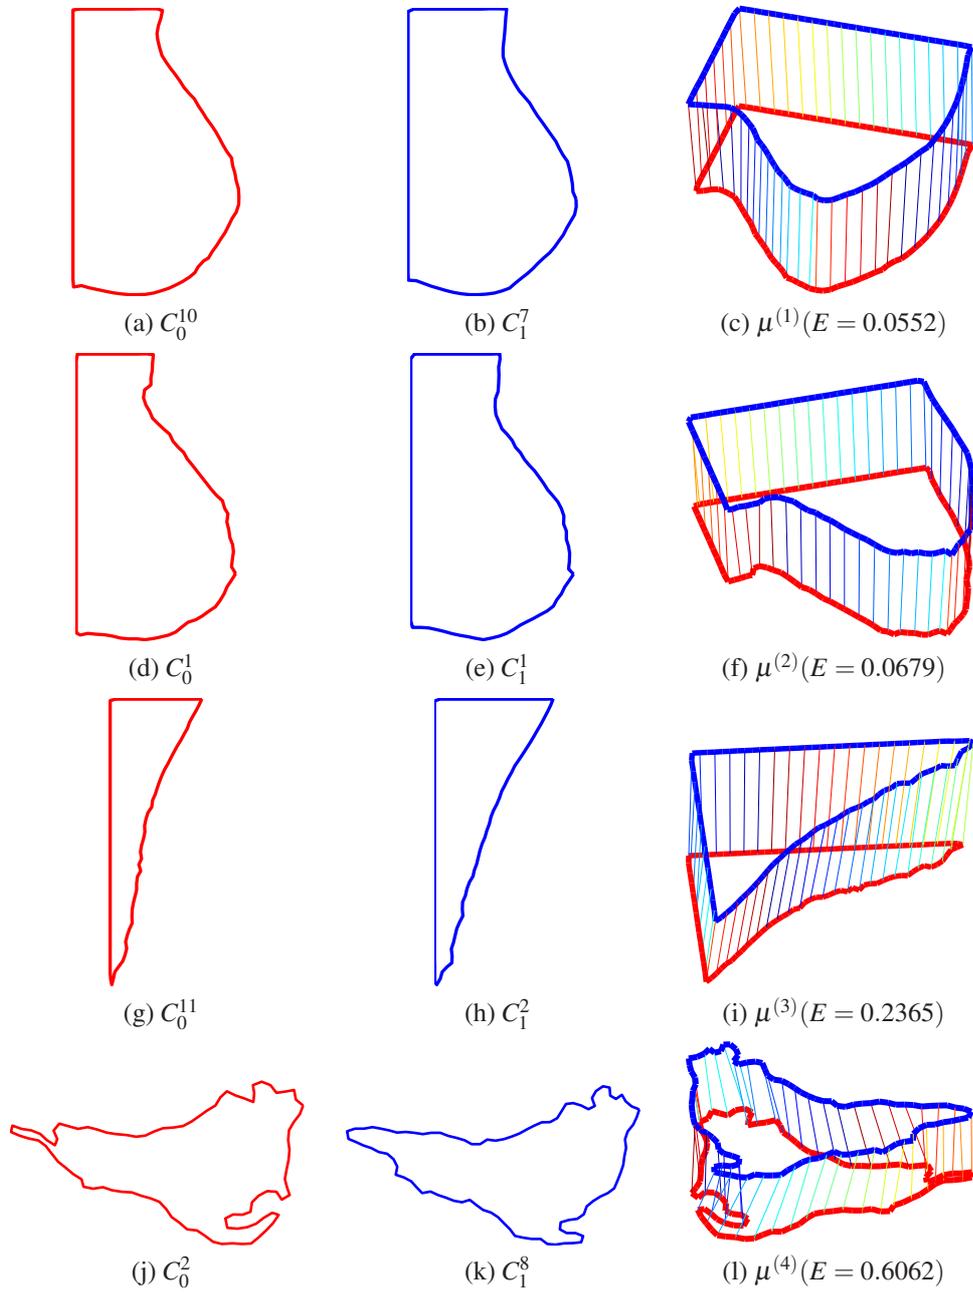


Figure 4: Correspondences between each pair of corresponding contours.

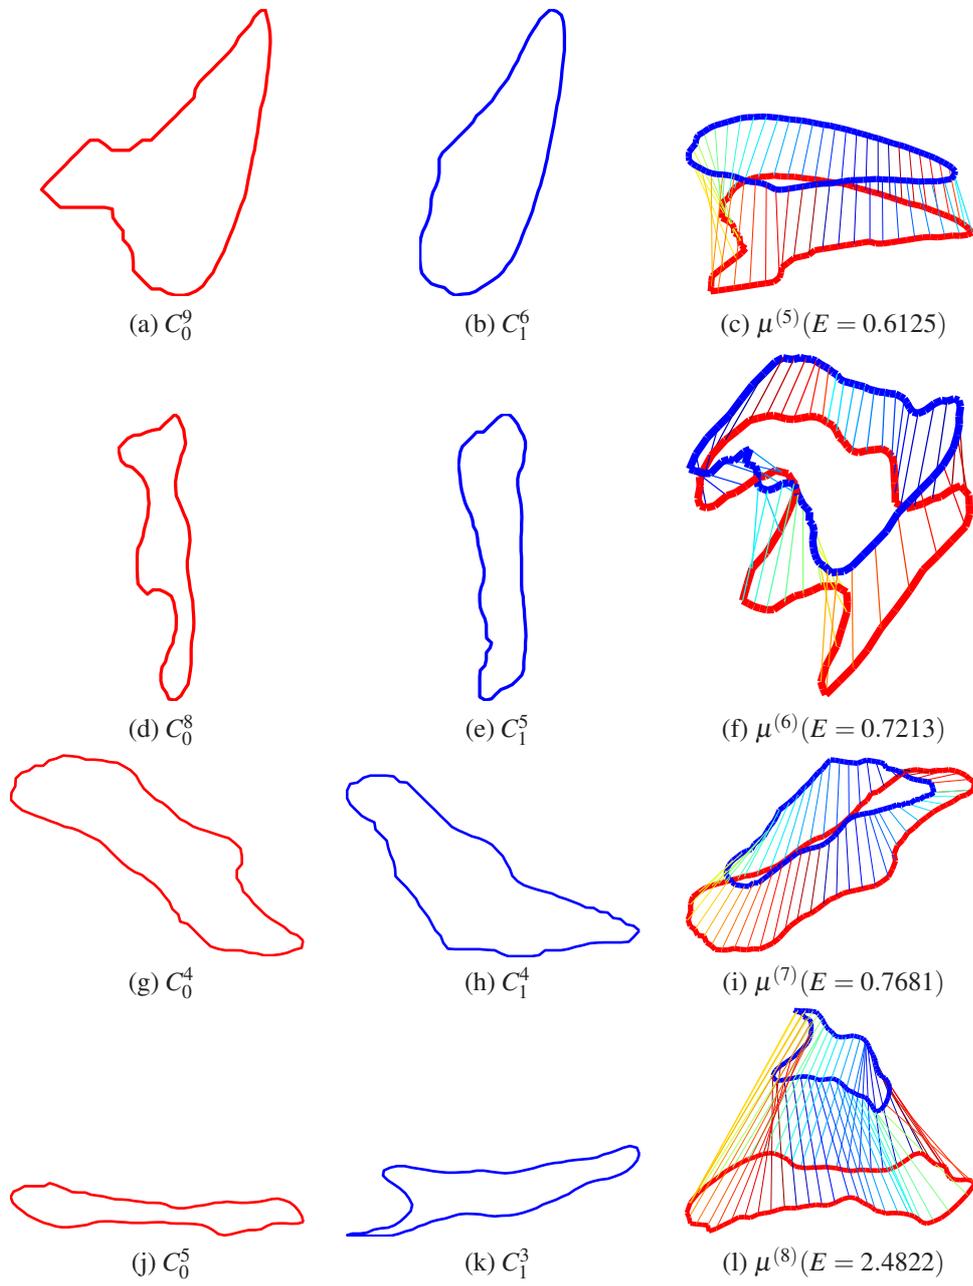


Figure 5: Correspondences between each pair of corresponding contours.

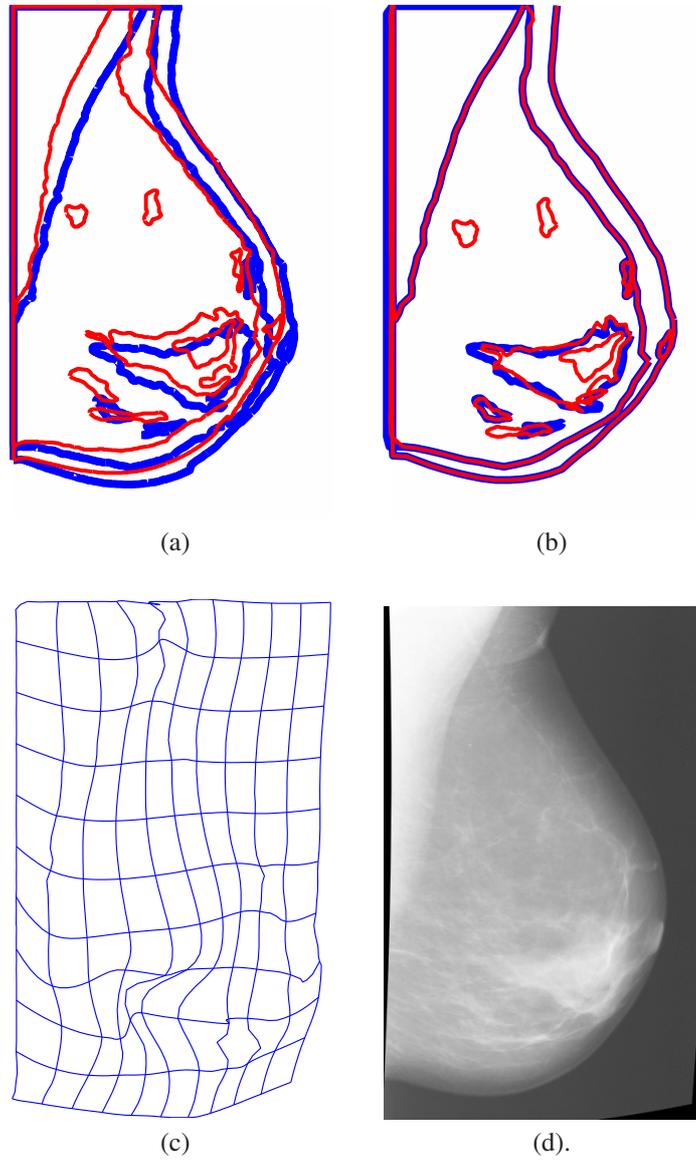


Figure 6: (a) Original $SC(I_0)$ and original $SC(I_1)$ are superimposed. (b) Transformed $SC(I_0)$ are aligned to original $SC(I_1)$. (c) Deformation grid of the optimal transformation. (d) Warped image of I_0 according to the optimal deformation.