

This is a preprint of material published in ISSN 0302-9743, ISBN 978-3-642-33414-6, Proceedings of the 15th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI 2012), Part I, LNCS 7510, pp. 438-446, 2012.

© Springer-Verlag Berlin Heidelberg 2012

Data-Driven Breast Decompression and Lesion Mapping from Digital Breast Tomosynthesis

M. Wels¹, B. M. Kelm¹, M. Hammon², A. Jerebko³, M. Sühling¹, and D. Comaniciu⁴

¹ Siemens AG, Corporate Technology, Erlangen, Germany

² University Hospital Erlangen, Department of Radiology, Germany

³ Siemens AG, Healthcare, Erlangen, Germany

⁴ Siemens Corporation, Corporate Research and Technology, Princeton, NJ, USA

Abstract. Digital Breast Tomosynthesis (DBT) emerges as a new 3D modality for breast cancer screening and diagnosis. Like in conventional 2D mammography the breast is scanned in a compressed state. For orientation during surgical planning, e.g., during presurgical ultrasound-guided anchor-wire marking, as well as for improving communication between radiologists and surgeons it is desirable to estimate an uncompressed model of the acquired breast along with a spatial mapping that allows localizing lesions marked in DBT in the uncompressed model. We therefore propose a method for 3D breast decompression and associated lesion mapping from 3D DBT data. The method is entirely data-driven and employs machine learning methods to predict the shape of the uncompressed breast from a DBT input volume. For this purpose a shape space has been constructed from manually annotated uncompressed breast surfaces and shape parameters are predicted by multiple multi-variate Random Forest regression. By exploiting point correspondences between the compressed and uncompressed breasts, lesions identified in DBT can be mapped to approximately corresponding locations in the uncompressed breast model. To this end, a thin-plate spline mapping is employed. Our method features a novel completely data-driven approach to breast shape prediction that does not necessitate prior knowledge about biomechanical properties and parameters of the breast tissue. Instead, a particular deformation behavior (decompression) is learned from annotated shape pairs, compressed and uncompressed, which are obtained from DBT and magnetic resonance image volumes, respectively. On average, shape prediction takes 26 s and achieves a surface distance of 15.80 ± 4.70 mm. The mean localization error for lesion mapping is 22.48 ± 8.67 mm.

1 Introduction

According to the World Cancer Report 2008 (globocan.iarc.fr, 2012/01/23) breast cancer is the most frequent cancer diagnosis in women among all specifiable kinds of cancer. Early detection is assumed to significantly improve outcomes. That is why breast cancer screening is recommended by many national organizations for most older women, e.g., the American Cancer Society (www.cancer.org,




Fig. 1: The MLO image acquisition (a), a 3D DBT scan (b), and a 3D rendering of an uncompressed breast with mapped lesions (c). Radiographic image with courtesy of University Hospital Erlangen.

2012/05/31) recommends yearly screening mammography for women aged 40 and above as long as they are in good health.

Recently, Digital Breast Tomosynthesis (DBT) increasingly replaces common 2D mammography for differential diagnosis and is in discussion for screening [1]. It provides 3D image volumes of the compressed breast (see Fig. 1(b)) that are reconstructed from multiple 2D projections acquired at varying angles. Being a 3D imaging modality DBT naturally allows superior spatial localization of suspicious lesions. For a mediolateral-oblique (MLO) scan the breast is compressed as sketched in Fig. 1(a). Typically, a second scan is acquired in craniocaudal (CC) direction during an examination. For surgical planning it is common clinical practice to mark the lesions in the scans and to communicate the rough localization of suspicious findings in the uncompressed breast via schematic 2D drawings. The latter naturally suffers from inaccuracies and can often only be dissolved by additional, potentially ionizing and costly, imaging. Providing more accurate lesion localization in the uncompressed breast, e.g., in terms of a 3D rendering view (see Fig. 1(c)), without additional imaging has the potential to facilitate surgical planning and related procedures, e.g., placing pre-operative markers, at low cost.

We therefore propose a method for estimating uncompressed 3D breast shapes from 3D DBT MLO scans. The reconstructed uncompressed shape and the original compressed shape depicted in the DBT scan establish a reference frame that can be used to map lesions found in the image of the compressed breast to the corresponding location in the uncompressed breast. Essentially, our method consists of two major steps: shape prediction and lesion mapping. For shape prediction the input data is transformed to a representation suitable for multiple multi-variate regression. This is done by fully-automatically detecting the nipple (or papilla) from the DBT images, segmenting the breast area, extracting the breast surface, and canonically re-sampling it. For lesion mapping we make use of the point correspondences established by applying the same canonical re-sampling to the uncompressed shapes used for predictor training and shape

model construction. This allows us to compute a thin-plate spline (TPS) interpolation for lesion mapping.

Despite the huge body of literature on biomechanical breast modeling [2, 3] and its applications in the context of 2D mammography [4] there are few publications dealing with DBT. As a notable exception van Schie et al. [5] match corresponding regions from ipsilateral DBT views (MLO and CC). Unlike in our case, the behavior of breasts—that are assumed to be hemispheres—under compression/decompression is explicitly modeled by approximating breast tissue properties. Regions are mapped by intermediately mapping them to a decompressed version of the initial geometric model that has been matched to the compressed breast before. The matching region in the ipsilateral view is finally found after rotation and repeated compression. This is different from our approach as we predict the shape of the uncompressed breast directly without intermediate steps and without explicitly modeling tissue behavior. We rather rely on a purely data-driven approach inherently capturing tissue behavior as it is present in the available training data.

2 Methods

2.1 Input Data and Feature Extraction

For shape prediction feature vectors $\mathbf{x} \in \mathbb{R}^K$ of uniform length K have to be extracted from the 3D DBT image volumes. For this we first segment the breast tissue area by thresholding and region growing, then extract the breast surface using Marching Cubes [6], and canonically re-sample it starting from the nipple to a fixed number of surface points. This re-sampling scheme is also relevant to lesion mapping, which will be discussed in more detail later. Finally, a high-dimensional feature vector suitable for machine learning-based shape prediction is obtained by composing the distances between every individual surface point and the nipple. Here, the nipple position is automatically determined by a machine learning-based landmark detection algorithm using 3D Haar-like features [7, 8]. It has been rapidly prototyped using an Integrated Detection Network (IDN) [9]. Note that, for processing, any breast is treated as a left breast, i.e., right breast images are mirrored.

2.2 Target Shape Model and 3D Shape Reconstruction

For the purpose of shape prediction we construct a statistical shape model [10] of the target shapes, i.e., the uncompressed breasts, from real patient data for the following two reasons. First, we are interested in a realistic but smooth estimate of the uncompressed breast shape omitting artefactual shape variations of individual examples. Second, the statistical shape model significantly reduces the number of shape parameters to be estimated (see below). From a statistical point of view, this simplifies the prediction problem, allowing successful training with less training data. Thus we strive for capturing as much shape information as possible in as few parameters as possible.

Here, we employ a simple linear shape model [10]: a shape $\mathbf{s} = (x_n, y_n, z_n)_{n=1, \dots, N}$ consisting of N points is described as linear combination $\mathbf{s} = \bar{\mathbf{s}} + \mathbf{P}\mathbf{y}$ where \mathbf{P} is a matrix formed by base vectors of the shape space, $\bar{\mathbf{s}}$ is the mean shape, and $\mathbf{y} = (y_1, \dots, y_L)$ are the shape parameters. The base vectors are found by Principal Component Analysis (see reference [10]). The amount of variation captured by the model is determined by the number of principal components L included in the model. Here, it turns out that

only $L = 5$ shape parameters are sufficient for preserving 90% of the total variance in a sample of 74 breasts (see Fig. 2). After verifying the shape space projections of these 74 breasts and given our training base for shape prediction we consider 5 shape parameters to be sufficient for our current system. This restriction also leverages the problem of the imperfect ground-truth at hand, i.e., larger uncompressed breasts within our data collection are occasionally deformed by the coils used for MR scanning as depicted in Fig. 4 (c). Note that in order to align the model shapes and establish point correspondence we re-sample relative to the same anatomical entities as for the compressed DBT shapes.




Fig. 2: The fraction of total shape variance plotted against the number of used shape parameters.

2.3 3D Shape Prediction by Multiple Multi-Variate Random Forest Regression

Predicting $L = 5$ real-valued shape coefficients from a K -dimensional feature vector can be cast as the multiple multi-variate regression problem $f : \mathcal{X} \rightarrow \mathcal{Y}$ with $\mathbf{y} = f(\mathbf{x})$, $\mathbf{x} \in \mathcal{X} = \mathbb{R}^K$, $\mathbf{y} \in \mathcal{Y} = \mathbb{R}^L$.

Here, we have chosen to use a Random Forest (RF) regressor [11] since it can capture linear as well as nonlinear dependencies and is known to achieve good generalization performance in general. While the more popular RF classifier employs randomized *decision* trees, RF regression uses an ensemble of M randomized *regression* trees. Like classification trees, regression trees partition the domain \mathcal{Y} of response variables \mathbf{y} into J regions $\mathcal{R} = \{\mathcal{R}_j | j = 1, \dots, J\}$ with $\mathcal{Y} = \bigcup_{j=1}^J \mathcal{R}_j$ and $\bigcap_{j=1}^J \mathcal{R}_j = \emptyset$. Unlike classification trees, regression trees model the distribution of response variables for any region R_j as a joint normal distribution with constant mean $\boldsymbol{\mu}_j$ and covariance $\boldsymbol{\Sigma}_j$. Since the response variables \mathbf{y} are the shape parameters of a linear statistical shape model here, they are uncorrelated by construction [10]. Hence, we assume that $\boldsymbol{\Sigma}_j = \text{diag}(\sigma_{j_1}^2, \sigma_{j_2}^2, \dots, \sigma_{j_L}^2)$ is a diagonal matrix with determinant $|\boldsymbol{\Sigma}_j| = \prod_{l=1}^L \sigma_{j_l}^2$. Randomized regression tree induction can be summarized as follows: given a training sample $\mathcal{T} = \{(\mathbf{x}_i, \mathbf{y}_i) | i = 1, \dots, I\} \subset \mathcal{X} \times \mathcal{Y}$ with cardinality T , regions \mathcal{R}_ν are recursively

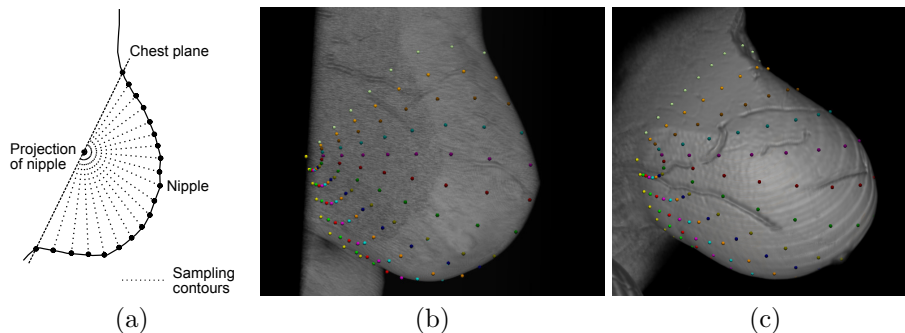


Fig. 3: The anatomical entities and the contours of the surface re-sampling planes in the MLO plane (a), and the canonical surface re-sampling scheme applied to a DBT scan (b) and an MR scan (c) of a compressed and uncompressed breast, respectively. Radiographic images with courtesy of University Hospital Erlangen.

split into $\mathcal{R}_\nu = \mathcal{R}_{\nu_1} \cup \mathcal{R}_{\nu_2}$ with $\mathcal{R}_{\nu_1} \cap \mathcal{R}_{\nu_2} = \emptyset$, where each region \mathcal{R}_ν is associated with a node ν of a binary tree and a subset $\mathcal{T}_\nu \subset \mathcal{T}$ of the training sample with cardinality T_ν . Each split is defined by a threshold t and an index m to one of the features. It partitions \mathcal{T}_ν into subsets $\mathcal{T}_{\nu_1} = \{(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{T}_\nu | x_{i_m} \leq t\}$ and $\mathcal{T}_{\nu_2} = \{(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{T}_\nu | x_{i_m} > t\}$. The optimal split is chosen among randomly selected features as to maximize the decrease in sample entropy

$$\Delta i(\nu) = T_\nu \cdot \ln |\Sigma_\nu| - (T_{\nu_1} \cdot \ln |\Sigma_{\nu_1}| + T_{\nu_2} \cdot \ln |\Sigma_{\nu_2}|).$$

During application, i.e., shape prediction, a feature vector \mathbf{x} is passed to appropriate leaf nodes ν_m according to the induced feature thresholds for each regression tree $m \in \{1, \dots, M\}$. Finally, the mean values $\boldsymbol{\mu}_j = \boldsymbol{\mu}_{\nu_m} \in \mathcal{Y}$ from the leaf nodes (corresponding to a region R_j) are averaged for prediction.

2.4 Thin-Plate Spline-Based Lesion Mapping from Corresponding Surface Points

A TPS mapping [12], which is computed based on two input shapes \mathbf{s} and \mathbf{s}' , is a function $f_{TPS} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ that maps each point of the input shape to its corresponding point in the target shape, i.e., $f_{TPS}(x_n, y_n, z_n) = (x'_n, y'_n, z'_n)$. Its parameters are based on closed-form solutions minimizing the integral quadratic variation for each coordinate axis. Being a continuous mapping it interpolates between the known data points and can therefore be used for mapping lesions $\mathbf{p} = (x, y, z)^T$ from within the compressed breast \mathbf{s} to within the predicted uncompressed breast $\mathbf{s}' = \bar{\mathbf{s}} + \mathbf{P}\mathbf{y}$.

For that reason we establish point correspondences between compressed and uncompressed shapes \mathbf{s} and \mathbf{s}' . Both the source shapes and the target shapes are parameterized relative to common anatomical entities that are identifiable in both cases. Figure 3 (a) summarizes the scheme with respect to an MLO scan.

- The first entity used is the nipple. It is automatically detected in the DBT images and has been manually annotated in the uncompressed breast shapes.




Fig. 4: Exemplary results for shape prediction on unseen data (a–e). The contour of the predicted shape is depicted in green, the one of the ground-truth annotation in red. Radiographic images with courtesy of University Hospital Erlangen.

- The second entity is the chest plane, which we define as the plane orthogonal to the MLO plane intersecting with the inframammary fold and its counterpart at the top of the DBT MLO scan. Both points are detected in the image slice aligning with the MLO plane and going through the nipple of the 3D MLO scan by analyzing the surface contour for curvature characteristics. In the uncompressed breast shapes these points have been manually annotated.
- The rotation axis for surface re-sampling is the axis orthogonal to the MLO plane intersecting with the nipple’s projection onto the chest plane.
- Further surface points within the MLO plane through the nipple are equidistantly sampled between the nipple and the inframammary fold and between the nipple and the top delineation of the breast. Then, intersection contours between the surface and the planes rotated around the rotation axis defined before and passing through the re-sampled points in the MLO plane are extracted. Equidistant re-sampling of these contours yields the remaining points of the re-sampled surface. This is done the same way for surfaces originating from DBT scans as well as uncompressed breast surfaces.

Figure 3 (b+c) shows a compressed/uncompressed breast pair where the described canonical re-sampling scheme has been applied.

3 Experimental Setting and Results

Nipple detection has been trained and evaluated on 122 annotated training data sets, i.e., DBT scans. The 8-fold cross-validated mean position error is 5.89 ± 6.56 mm, which reflects the annotation uncertainty.

The shape model has been computed from a population of 74 surface breast shapes. They all have been manually annotated in MR data showing hanging breasts as the patients all lay prone at the time of image acquisition. Evaluation of shape prediction has been carried out on 24 breast shape pairs (compressed/uncompressed). For these 24 shapes, DBT images and associate input shapes as well as ground-truth uncompressed target shapes from MR are available as the patients underwent both types of examination. However, none of the MR scans of these 24 patients has been used for shape model generation. For evaluation of lesion mapping, an expert radiologist has annotated 49 correspond-

ing points inside the breast both in the DBT scan and in the accompanying MR scan. These expert annotations are used to assess accuracy of lesion mapping.

The projection of the original breast shapes into the 5-dimensional shape space yields nice and smooth shapes but also a mean average surface distance of already 10.57 ± 4.36 mm. Using fully-automatic nipple detection and the proposed regression approach, the mean average surface distance slightly increases to 15.80 ± 4.70 mm (result from 8-fold cross-validation). Exemplary results are depicted in Fig. 4, which also demonstrates the desired smoothing behavior of our method (see Fig. 4(c)). The evaluation of the final lesion mapping step yields a mean localization error of 22.48 ± 8.67 mm (see Figs. 5 and 6). Note, however, that this error also includes the intentional discrepancy between the smoothed breast surface after shape space projection and the annotated breast surface in MR. Thus, it cannot be compared with accuracies obtained by methods that have access to the exact target shape. Feature extraction takes 26 s on average where 86% of the time is spent detecting the nipple. Shape prediction takes 36 ms. Comparison of our results to existing methods for lesion mapping is difficult due to the lack of competing methods addressing exactly the same problem. In spite of this, van Schie et al. [5] report position errors for lesion mapping between ipsilateral DBT scans in a comparable range, which they consider reasonably accurate. Further improvement is certainly desirable for clinical application of our method.




Fig. 5: The histogram of localization errors achieved for lesion mapping.

4 Conclusions

In this paper we have presented a novel method for uncompressed breast shape prediction from 3D DBT image volumes. As a core contribution, the method does not involve any explicit biomechanical modeling of deformations, e.g., induced by gravity or compression. Instead, we make use of data-driven multiple multivariate RF regression to immediately predict uncompressed target shapes in terms of their shape parameters describing them in a statistical shape space. We apply a TPS mapping to map lesions identified in the input data to the corresponding location in the predicted uncompressed breast. For that, breast surface point correspondences have been defined relative to anatomical entities that are identifiable both in the compressed and in the uncompressed breast. These entities are fully-automatically detected in the DBT scans. Processing one data set takes less than half a minute where most of the time is due to




Fig. 6: Lesion mapping results. The images show ground-truth annotations (a, c, and e) and associated locations estimated by our method (b, d, and f). Radiographic images with courtesy of University Hospital Erlangen.

feature extraction. The achieved lesion mapping accuracy is encouraging and supports further refinement of the proposed data-driven approach.

References

1. Gilbert, F.J., Young, K.C., Astley, S.M., Whelehan, P., Gillan, M.G.C.: Digital Breast Tomosynthesis. NHS Cancer Screening Programmes, Sheffield (2010)
2. Tanner, C., White, M., Guarino, S., Hall-Craggs, M.A., Douek, M., Hawkes, D.J.: Large breast compressions: Observations and evaluation of simulations. *Med. Phys.* **38**(2), 682–690 (2011)
3. Samani, A., Bishop, J., Yaffe, M.J., Plewes, D.B.: Biomechanical 3-D finite element modeling of the human breast using MRI data. *IEEE Trans. Med. Imag.* **20**(4), 271–279 (2001)
4. Zhang, Y., Qiu, Y., Goldgof, D.B., Sarkar, S., Li, L.: 3D finite element modeling of nonrigid breast deformation for feature registration in X-ray and MR images. In: *IEEE Workshop Appl. Comp. Vis.*, Austin, TX, USA, p. 38. (2007)
5. van Schie, G., Tanner, C., Snoeren, P., Samulski, M., Leifland, K., Wallis, M.G., Karssemeijer, N.: Correlating locations in ipsilateral breast tomosynthesis views using an analytical hemispherical compression model. *Phys. Med. Biol.* **56**(15), 4715–4730 (2011)
6. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. *Comput. Graph.* **21**(4), 163–169 (1987)
7. Zheng, Y., Barbu, A., Georgescu, B., Scheuering, M., Comaniciu, D.: Four-chamber heart modeling and automatic segmentation for 3D cardiac CT volumes using Marginal Space Learning and steerable features. *IEEE Trans. Med. Imag.* **27**(11), 1668–1681 (2008)
8. Wels, M., Zheng, Y., Carneiro, G., Huber, M., Hornegger, J., Comaniciu, D.: Fast and robust 3-D MRI brain structure segmentation. In: Yang, G.Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) *MICCAI 2009, Part II*. LNCS, vol. 5762, pp. 575–583. Springer Berlin/Heidelberg (2009)
9. Sofka, M., Ralovich, K., Zhang, N.B.J., Zhou, S.K.: Integrated Detection Network (IDN) for pose and boundary estimation in medical images. In: *IEEE Int. Symp. Biomed. Imag.: Nano To Macro*, Chicago, IL, USA, pp. 294–299 (2011)
10. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active Shape Models—their training and application. *Comput. Vis. Image Understand.* **61**(1), 38–59 (1995)
11. Breiman, L.: Random Forests. *Mach. Learn.* **45**(1), 5–32 (2001)
12. Bookstein, F.L.: Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell.* **11**(6), 567–585 (1989)