Mammogram Compression Using Super-Resolution

Jun Zheng[1], Olac Fuentes[1], Ming-Ying Leung[1], and Elais Jackson[2] *

[1] The University of Texas at El Paso, El Paso, Texas, 79968
jzheng@miners.utep.edu, ofuentes@utep.edu, mleung@utep.edu
[2] Tennessee State University, Nashville, TN 37209

Abstract. As mammography moves towards completely digital and produces prohibitive amounts of data, compression plays an increasingly important role. Although current lossless compression methods provide very high-quality images, their compression ratios are very low. On the other hand, several lossy compression methods provide very high compression ratios but come with considerable loss of quality. In this work, we describe a novel compression method that consists of downsampling the mammograms before applying the encoding procedure, and applying super-resolution techniques after the decoding procedure to recover the original resolution image. In our experiments, we examine the tradeoffs between compression ratio and image quality using this scheme, and show it provides significant improvements over conventional methods.

1 Background

As mammography moves towards completely digital, technological advances in data storage and transmission have not kept up with the tremendous growth of digital data. This creates serious challenges for long-term storage and efficient transmission of mammograms. For example, a typical mammogram can be 4500×4500 pixels. If stored in uncompressed 16-bit per pixel (bpp) format, it would take about 40 MB for storage and approximately half an hour for transmission using a high-speed modem [12]. Thus compression will play an increasingly important role in Picture Archiving and Communication Systems (PACS) to reduce file sizes while maintaining relevant diagnostic information.

In recent years, there has been discussion about which type of compression techniques, lossy or lossless, is better for mammogram compression. Although current lossless compression methods provide very high quality images, the compression ratios are very low, typically from 1.5:1 to 3:1. On the other hand, several lossy compression methods provide acceptable compression ratios but come with considerable loss of image quality and diagnostic information [3–8, 12].

In this work we describe a novel lossy compression method that consists of downsampling the mammograms before applying the encoding procedure, and

* This work was supported in part by NSF grants DMS0800272 and IIS0852066, and NIH grants 2506GM008012-39 and 2G12RR008124-16A1.

applying super-resolution techniques after the decoding procedure to recover the original resolution image.

2 Methods

Incorporating super-resolution into compression aims to maximize compression ratio while maintaining relatively high-quality images. The flowchart of the proposed method is shown in figure 1. The method first downsamples a high-resolution image $HR(x)$ to a low-resolution one $LR(x)$, then it uses an algorithm similar to JPEG to encode it. The decompression process first decodes the stored low resolution image to obtain $LR'(x)$ and then applies a super-resolution algorithm to produce $HR'(x)$.

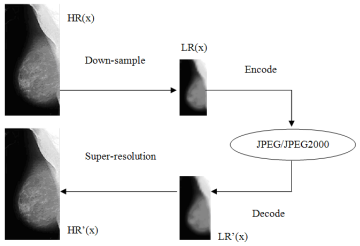

Fig. 1. flowchart of the method

In the pre-processing stage, we downsample the mammograms using bilinear downsampling. In the post-processing stage, the super-resolution algorithm generates high-resolution mammograms from low-resolution decoded mammograms with no manual registration. The super-resolution algorithm consists of four main steps. The first step consists of automatically aligning the breasts to a standardized position. The second step uses a process called eigentransformation to infer a global model representing the low-frequency information in the image. In eigentransformation, Principal Component Analysis (PCA) is used to fit the input images as a linear combination of the low resolution images in the training set. The HR images are then inferred by replacing the LR training images with HR ones, while retaining the same combination coefficients. In the third step, a patch-based one-pass algorithm generates the high-frequency contents of the HR images. The fourth step remaps the breasts back to their original position.

2.1 Automatic alignment

Image alignment is the key to the success of our automatic mammogram super-resolution algorithm. In practice, we cannot assume that any low-resolution

mammogram has been accurately aligned, even though the approximate positions of the breasts in mammograms are given by mammography sensors. Therefore, in preprocessing, we automatically align the mammograms to make sure each breast is exactly in the same position, and then perform the super-resolution. The automatic alignment process consists of two parts, segmentation-based initialization and 2-pass mesh warping [11].

The 2-pass mesh warping algorithm accepts a source image and two 2-D arrays of coordinates. The first array, S, specifies the coordinates of control points in the source image. The second array, D, specifies their corresponding positions in the destination image. The first pass is responsible for resampling each row independently. It maps all (u, v) points to their (x, v) coordinates in the intermediate image I. For each pixel P in intermediate image I, the value of P is evaluated as a weighted sum from the left most boundaries of P in S, x_0, and the rightmost boundaries of P in S, x_1.

$$P = \frac{\sum_{x=x_0}^{x_1} k_x S_x}{x_1 - x_0}$$

where k_x is the scale factor of source pixel S_x, and the subscript x denotes the index that lies between $floor(x_0)$ and $ceil(x_1)$. The scale factor k_x is defined as

$$k_x = \begin{cases} ceil(x) - x_0 & \text{if } floor(x) < x_0 \\ 1 & \text{if } x_0 \leq x < x_1 \\ x_1 - floor(x) & \text{if } ceil(x) > x_1 \end{cases}$$

The second pass then resamples each column in I, mapping every (x, v) point to its final (x, y) position. This process is virtually identical to the first pass; we just need to substitute (x, v) for (u, v), and substitute (x, y) for (x, v) [11].

The key to apply 2-pass mesh warping is to build the 2-D arrays of coordinates. The segmentation-based initialization builds the 2-D arrays of coordinates automatically. The initialization for 2-pass mesh warping consists of 5 steps (Figure 2):

1. Convert the input image to binary image.
2. Use erosion to remove the labels on the binary image.
3. Convert the binary image to a gradient image to find the edge.
4. Use skeletonization to reduce the edge to a single line.
5. Sample points on the line and make the 2D- array of coordinates for 2-pass mesh warping automatically.

2.2 Global modeling

In global modeling we use an algorithm called eigentransformation, which was originally introduced by Wang [9]. The eigentransformation is a simple and powerful technique for image enhancement based on principal component analysis

(a) (b) (c) (d) (e) (f)

Fig. 2. Automatic registration: (a) Original image. (b) Converting the original image to a binary image. (c) Using erosion to remove the labels in the binary image. (d) Converting the binary image to gradient image to find the edge. (e) Using skeletonization to reduce the edge to a single line. (f) Sampling points on the line and making the mesh.

(PCA). It assumes that we have a training set of pairs of images $\langle (L_1, H_1), \ldots, (L_n, H_n) \rangle$, where each pair (L_i, H_i) contains a low resolution image L_i and its corresponding high-resolution counterpart H_i. The eigentransformation allows any image to be represented as a linear combination of images in the training set. When given a low resolution image L, it finds the vector of coefficients $[c_1, \ldots, c_n]$ so that

$$L = \Sigma_{i=1}^{n} c_i L_i + \mu_L$$

where μ_L is the mean low-resolution images.

Given the vector $[c_1, \ldots, c_n]$, the approximate high resolution image H can be computed by

$$H = \Sigma_{i=1}^{n} c_i H_i + \mu_H$$

where μ_H is the mean high-resolution image.

Because the coefficients are not computed from HR training data, some noise-like distortion may be introduced. To reduce distortion, we apply constraints by bounding the projection onto each eigenvector by its corresponding eigenvalue, then the synthesized image is reconstructed from these constrained coefficients.

2.3 Local modeling

Given a global model, to construct the corresponding local model, we first filter the global model with a Gaussian high-pass filter, and then subdivide the filtered global model into patches, which we call the low-frequency patches of the high-resolution (HR) images, by scanning a window across the image in raster-scan order. Similarly, we also filter and subdivide the HR images in the training set into patches which we call high-frequency patches of the training HR images.

To construct a local model, for each low-frequency patch, a high-frequency patch of the training HR image is selected by a nearest neighbor search from the training set based on local low-frequency details and adjacent HR patches

previously determined. The selected high-frequency patch should not only come from a location in the training images that has a similar corresponding low-frequency appearance, but also agrees with the overlapping pixels, which we call high-frequency overlap, at the edges of its previously determined high-frequency neighbors. This ensures that the high-frequency patches are compatible with those of the neighboring high-frequency patches.

In this work we compute the local model with an algorithm similar to Freeman's [1][2] one-pass algorithm. We first concatenate the pixels in the low-frequency patch and the high-frequency overlap to form a search vector. The training set also contains a set of such vectors. Then we search for a match by finding the nearest neighbor in the training set. When we find a match we extract the corresponding high-frequency patch from the training data set and add it to the initial global model to obtain the output image.

Mathematically, this process can be described as follows. Suppose we have a training data set

$$\{(x^{(i,j,k)}, y^{(i,j,k)}, z^{(i,j,k)}),$$
$$i = 1, 2, \ldots, l; j = 1, 2, \ldots, m; k = 1, 2, \ldots, n\}$$

where $x^{(i,j,k)}$ is the low-frequency patch at the i^{th} row and j^{th} column of the k^{th} training HR image, $y^{(i,j,k)}$ is the corresponding high-frequency overlap and $z^{(i,j,k)}$ is the corresponding high-frequency patch of the training HR image, l is the number of rows of patches in a training image, m is the number of columns of patches in a training image and n is the number of training images.

Given an input LR patch $\overline{x}$, we need to find an HR patch $z^{(i',j',k')}$ such that

$$z^{(i',j',k')} = \min_{i,j,k}(d(\overline{x}, x^{(i,j,k)}) + \alpha * (d(y^{(i,j,k)}, y_N^{(i,j,k)})))$$

where $d(x, y)$ is the Euclidean distance between x and y, $y_N^{(i,j,k)}$ is the overlap of $z^{(i,j,k)}$ with the adjacent, previously determined high-frequency patches, which are the patches above and to the left of the current high-frequency patch in the local model, α is a user-controlled weighting factor, and $z^{(i',j',k')}$ is the selected high-frequency patch.

3 Experimental results

We use DDSM (Digital Database for Screening Mammography) for our experiments. DDSM is a standard dataset used by mammography image analysis research community. The database has about 2,500 cases. Each case includes two images of each breast, along with some associated patient and image information. In this work, we use 100 normal left Mediolateral Oblique (MLO) images from DDSM for training and 10 normal left MLO images images for testing.

In this paper, the quality of a super-resolution image is defined as the similarity of the decompressed image with the original high-resolution image. We use Peak Signal-to-Noise Ratio ($PSNR$) and Mean Structural Similarity ($MSSIM$) index to measure the quality of results.

Let X and Y be two images to be compared. The $PSNR$, which is most commonly used as a measure of quality of reconstruction, is defined as

$$PSNR(X,Y) = 20 \times log_{10} \frac{255}{RMSE(X,Y)}$$

where $RMSE$ is the root mean square error between the two images.

The structural similarity (SSIM) index [10] is an implementation of the idea of structural similarity, from an image formation point of view, which takes into account contrast, luminance, and structure to determine similarity between two images. $SSIM$ is defined as

$$SSIM(x,y) = \frac{(2\mu_x\mu_y + C1)(2\sigma_{xy} + C2)}{(\mu_x^2 + \mu_y^2 + C1)(\sigma_x^2 + \sigma_y^2 + C2)},$$

$$\sigma_{xy} = \frac{1}{T-1} \sum_{i=1}^{T} (x_i - \mu_x)(y_i - \mu_y)$$

where x and y are subimages of X and Y, T is the total number of pixels in each subimage, μ_x is the average of x, μ_y is the average of y, σ_x is the standard deviation of x, σ_y is the standard deviation of y. $C1 = (k_1 L)^2$ and $C2 = (k_2 L)^2$ are two variables to stabilize the division with small denominators, L is the dynamic range of the pixel values (typically this is 255), $k_1 = 0.01$ and $k_2 = 0.03$ by default. The mean SSIM (MSSIM) is then simply the mean of the SSIMs for all subimages. A value of MSSIM of 1 indicates perfect similarity [10].

Figures 3 and 4 show the image quality measures for our algorithm. The downsample factors in this experiment are 4, 8, 16, and the number of Discrete Cosine Transform (DCT) coefficients used in an 8×8 block for JPEG encoding are 1, 2, 3, 5, 10, and 15. Table 1 compares the results of our method (downsample factor = 16, 15 DCT coefficients) with the results of JPEG2000 and lossless JPEG. Sample results of different resolutions are reported in Figure 5.

Compression methods	PSNR	Compression ratio
JPEG2000	41.95	80:1
Lossless JPEG	lossless	3:1
Our method	35.9023	20480:1

Table 1. Average compression results for 10 mammograms

4 Discussion

From figures 3 and 4, we can see that when we just use 1 or 2 DCT coefficients, which results in very blocky images, we get a PSNR of 33.458 after super-resolution of this kind of blocky images. These results indicate that

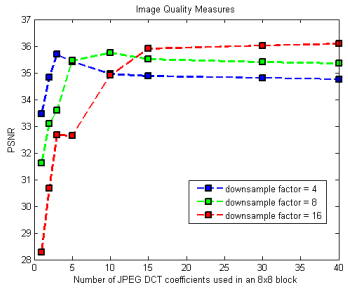

Fig. 3. PSNR of the decompressed images after super-resolution

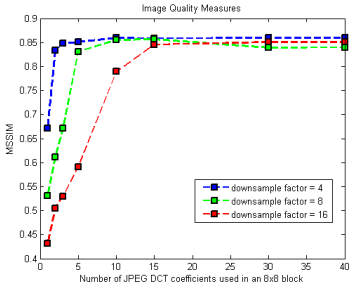

Fig. 4. MSSIM of the decompressed images after super-resolution

super-resolution can also attenuate the JPEG compression artifacts when the compression ratio is high. From Table 1, we can see that compared with the results of JPEG2000 and lossless JPEG, the compression ratio of our method is thousands of times higher than lossless JPEG, and hundreds of times higher than JPEG2000 with a slightly lower PSNR. However, some details of the decompressed images are different from the original images. Current work by our group consists of determining whether super-resolution affects clinical diagnostic performance, by both humans and computers.

For future work, we will study a hybrid scheme that is appropriate for accurate compression of mammograms. We will use lossless compression in the region of interest, and lossy compression with super-resolution in the other regions. This strategy has the potential to achieve high compression ratios without reducing the diagnostic quality of the mammograms.

References

1. Freeman, W.T., Jones, T.R., Pasztor, E.C.: Example-based super-resolution. IEEE Computer Graphics and Applications 22(2), 56–65 (2002)

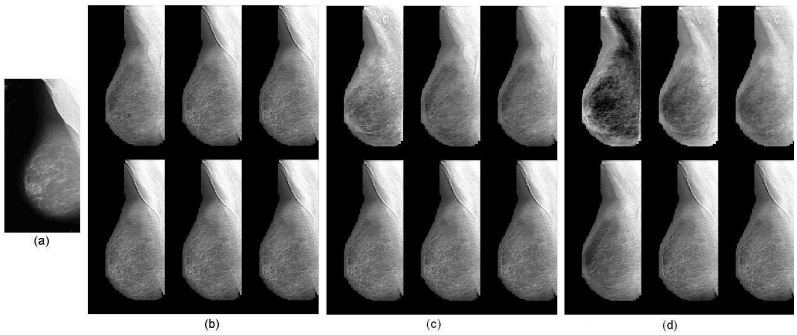

Fig. 5. (a) Original image. (b) The decompressed images after super-resolution (Coeffs=1,2,3,5,10,15; downsample factor=4).(c) The decompressed images after super-resolution (Coeffs=1,2,3,5,10,15; downsample factor=8).(d) The decompressed images after super-resolution (Coeffs=1,2,3,5,10,15; downsample factor=16).

2. Freeman, W.T., Pasztor, E.C., Carmichael, O.T.: Learning low-level vision. International Journal on Compter Vision 40(1), 25–47 (2000)
3. Grinstead, B., Sari-Sarraf, H., Mitra, S., Gleason, S.: Content-based compression of mammograms for telecommunication and archiving. 13th IEEE Symposium on Computer-Based Medical Systems (CBMS'00) 0, 37 (2000)
4. Li, X., Shen, Y., Ma, J.: An efficient medical image compression scheme. In: Proceedings of the 2005 IEEE Engineering in Medicine and Biology 27th Annual Conference. pp. 3437–3439. Shanghai, China (January 2006)
5. M.S., R.K., Koliwad, S., G.S., D.: Lossless compression of digital mammography using fixed block segmentation and pixel grouping. In: Proceedings of indian conference on computer vision, graphics and image processing. pp. 201–206 (2008)
6. Perlmutter, S.M., Cosman, P.C., Gray, R.M., Olshen, R.A., Ikeda, D., Adams, C.N., Betts, B.J., Williams, M.B., Perlmutter, K.O., Aiyer, J.L., Fajardo, L., Birdwell, R., Daniel, B.L.: Image quality in lossy compressed digital mammograms. Signal processing 59(2), 189–210 (June 1997)
7. Persons, K., Palisson, P., Manduca, A., Erickson, B.J., Savcenko, V.: An analytical look at the effects of compression on medical images. Journal of Digital Imaging 10, 60–66 (August 1997)
8. Scharcanski, J.: Lossless and near-lossless compression for mammographic digital images. IEEE International Conference on Image Processing pp. 2253–2256 (2006)
9. Wang, X., Tang, X.: Hallucinating face by eigentransformation. IEEE Transactions on Systems, Man and Cybernetics, Part C: Applications and Reviews (2005)
10. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. IEEE Transactions on Image Processing 13(4), 600–612 (April 2004)
11. Wolberg, G.: Digital Image Warping, chap. 2-Pass Mesh Warping, pp. 222–240. IEEE Computer Society Press (1990)
12. Zukoski, M.J., Boult, T., Iyriboz, T.: A novel approach to medical image compression. International Journal of Bioinformatics Research and Applications (2006)