Efficient Registration of Aliased X-Ray Images

M. Dirk Robinson[*] , Sina Farsiu[†], Joseph Y. Lo[‡], Peyman Milanfar[§], Cynthia A. Toth[¶]

[*]Ricoh Innovations, Menlo Park, CA, Email: dirkr@rii.ricoh.com
[†] Dept. of Ophthalmology, Duke University Medical Center, Durham, NC, Email: sina.farsiu@duke.edu
[‡]Dept. of Radiology and Biomed. Engr., Duke University Medical Center, Durham, NC, Email: joseph.lo@duke.edu
[§]Dept. of Electrical Engr., University of California, Santa Cruz, CA, Email: milanfar@ee.ucsc.edu
[¶] Dept. of Ophthalmology and Biomed. Engr., Duke University Medical Center, Durham, NC, Email: cynthia.toth@duke.edu

Abstract— **Multiframe image reconstruction produces images beyond the native resolution of a digital image sensor by way of accurate sub-pixel registration of aliased images. We present a novel multiframe registration approach for the purpose of enhancing resolution of digital mammogram images. We demonstrate the ability to improve resolution while maintaining normal radiation dosages.**

I. Introduction

While relatively new, digital mammography is rapidly replacing film-based mammography for the screening and diagnosis of early carcinomas in women. Solid-state detectors have demonstrated improved performance in terms of specificity and sensitivity over film-based imaging for certain groups of women such as those with dense breast tissue, women under the age of fifty, and pre-menopausal women [1]. Unlike film-based mammography, digital mammography provides the opportunity to apply sophisticated digital processing techniques to aid in screening.

The goal of designing digital mammography systems is to expose the patient to the minimum amount of radiation required to accomplish the screening task. Digital mammography systems face the same design tradeoff between image resolution, signal to noise ratio (SNR), and illumination or radiation strength as those found in any digital imaging system. Shrinking the pixel dimension at the detector increases sampling resolution at the expense of dynamic range and SNR. While the SNR and dynamic range may be recovered through multiple images, the increase in the total radiation exposure is often unacceptable. Alternatively, using large detector pixels improves both the dynamic range and the SNR of the system at the obvious expense of resolution. This approach is often avoided due to the resolution required to maintain the sensitivity of the mammography screen.

We propose fusing multiple low-dosage images, each containing spatial shifts, to produce a single high quality image at a resolution greater than the sampling rate of the detector. Such processing, commonly called super-resolution, has received much attention in recent years in the image processing community. We refer the interested reader to [2], [3] for a broad review of recent algorithmic development in this area (note that improving the resolution of X-Ray images is also discussed in [4], where the presented super-resolved X-Ray images in that paper were provided by the authors of this paper). Perhaps the most fundamental component to multiframe super-resolution is the accurate registration of aliased images.

Digital X-ray imaging systems often contain aliasing artifacts. It is these aliasing artifacts which enable multi-frame super-resolution processing to extract information beyond the sampling rate of the detector. Multiple captured images with phase variations allow super-resolution algorithms to *unwrap* these aliasing artifacts, reconstructing the high resolution image provided that accurate estimates of the subpixel motion can be obtained. Unfortunately, motion estimation of aliased images is an ill-posed problem and specially difficult in the low-SNR cases [5]. Inaccurate motion estimation is one of the main reasons why many super-resolution algorithms, which produce excellent results when applied to simulated data, fail to produce super-resolved images when applied to real image sequences. To overcome this problem, in our previous publication [6], we treated the motion estimation errors as outliers and relied on robust statistics to remove such artifacts from the final super-resolved image. While the noted robust technique has proven to be successful for variety of applications, in some medical applications it might not provide optimal solution. Of course, if accurate subpixel motion estimation is available, less information is discarded as outliers. As proven in [5], to achieve accurate subpixel motion estimation for aliased images, instead of common pairwise image registration schemes, multi-frame schemes should be exploited [7], [8].

In this paper, we describe a statistically optimal and numerically efficient technique for registering *sets* of aliased low dosage X-ray images. The results show improved resolution beyond the resolution of the detector offering system designers new tradeoffs in system design. In Sect. II, we describe the challenges associated with multiframe reconstruction of digital mammogram images. We present a novel two-stage algorithm for jointly registering and reconstructing a set of aliased, low-resolution images. Sections II-A and II-B describe the multiframe block matching stage and the multiframe subpixel registration and reconstruction stage, respectively. Section III presents the experimental results and concluding remarks are given in Section IV.

II. MULTIFRAME REGISTRATION AND RECONSTRUCTION

Application of standard super-resolution processing to the problem of digital mammography faces several significant challenges. For example, perfect immobilization of both the patient and the imaging system is nearly impossible during multiple acquisitions. As a results, the relative motion between the captured images rarely follows the simple translational motion model commonly assumed by super-resolution algorithms. When departing from the translational motion model, super-resolution processing becomes significantly more computationally demanding. Adding to the computational complexity is the size of the capture images. Our experimental system produces 10 megapixel 16 bit images (20 megabytes per image). Combining sets of such large images using floating point computations is numerically infeasible for typical desktop processors due to the memory limitations. Finally, the captured data has extremely low SNR. To minimize total exposure, we must use very low dosages of illuminating radiation. For example, Fig. 1 shows the computed SNR[1] as a function of exposure dosage (mAs) for a our phantom breast object. Our multiframe imaging approach is targeting very low exposure levels of around 10 mAs corresponding to roughly 3-4 db. Standard registration algorithms perform poorly at such low
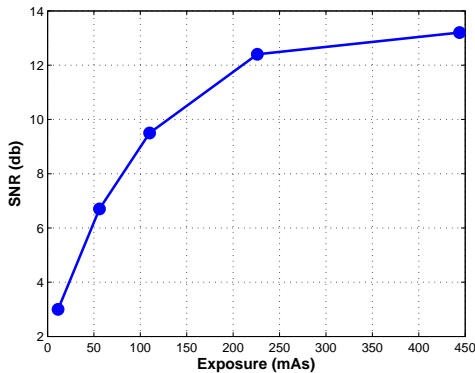

Fig. 1. The graph shows experimentally observed signal to noise ratios (SNR) for different exposure levels (mAs). The SNR was computed using images of a phantom breast resolution chart.

SNR limiting the ability to perform enhance the resolution.

We address the computational challenge by applying a novel super-resolution approach in a tile-based fashion. When using tile-based processing, the relative shifts between data sets is well approximated by the translational (or rotational) motion model. Furthermore, applying the algorithm to small tiles reduces the memory requirements of the multiframe reconstruction algorithm. Finally, the small tiles allow us to perform more sophisticated multiframe registration algorithms which *jointly* estimate the set of unknown registration parameters.

In our forward model, we assume that the observed images $\mathbf{y}_k$ are composed of translated $B \times B$ pixel blocks or tiles of

[1]In this work, the SNR was computed numerically as SNR= $20 log_{10} \frac{s}{n}$, where s is half the pixel value between the black and white signal regions and n is the noise standard deviation in flat regions.

the downsampled high resolution image $\mathbf{x}$. We represent this as

$$\mathbf{y}_k = \sum_t \mathbf{R}_t(\mathbf{o}_k)\mathbf{DS}(\mathbf{m}_{k,t})\mathbf{z}_t + \mathbf{e}_k \qquad (1)$$

where the vector lexicographically ordered $\mathbf{y}_k$ represents $N_L \times N_L$ samples of the observed image undersampled with respect to the high resolution image $\mathbf{z}$ by a factor of d in each dimension. The vector $\mathbf{z}_t$ represents samples of the tth unknown $dB \times dB$ high resolution image tile lexicographically ordered as a $(dB)^2 \times 1$ vector ($\mathbf{z}_t$ is the blurred version of $\mathbf{x}_t$ [6] and the deblurring details are omitted in this paper due to space limitations). The warping operator $\mathbf{S}(\mathbf{m}_{k,t})$ of size $(dB)^2 \times (dB)^2$ represents the sub-pixel spatial shifts between similar tiles in the captured images. The spatial shifting is described by the vector $\mathbf{m}_{k,t}$ for the kth frame. In our model, we assume that these spatial shifts are continuous values in the range of $[-d, d]$. This corresponds to the range of sub-pixel motions in the captured images. The downsampling operator $\mathbf{D}$ of size $B^2 \times (dB)^2$ captures the undersampling of the detector. We assume that the captured images are undersampled by an integer factor d. The $N_L \times B^2$ matrix $\mathbf{R}_t(\mathbf{o}_k)$ represents the super-pixel shift of the downsampled block in the captured image coordinates. The block is shifted by $\mathbf{o}$ pixels. Finally, $\mathbf{e}_k$ of size $N_L^2 \times 1$ represents the noise inherent in the analog-to-digital conversion. For our purposes, we assume this noise to be uncorrelated zero-mean noise with standard deviation σ.

The general problem of super-resolution is to combine $K + 1$ captured low-resolution images and estimate the high resolution image $\mathbf{z}$. To accomplish this task, the algorithm must also estimated the collection of unknown motion vectors. Without loss of generality, we assume that the initial image $\mathbf{y}_0$ defines the coordinate system of the high resolution image and hence we only have to estimate the unknown motion parameters for K images for a given set of $K + 1$ low resolution frames. To simplify the presentation, we use an underline notation to represent the larger set of unknown motion vectors $\underline{\mathbf{o}} = [\mathbf{o}_1, \ldots, \mathbf{o}_K]^T$ and $\underline{\mathbf{m}} = [\mathbf{m}_1, \ldots, \mathbf{m}_K]^T$.

The algorithm follows a two-stage process to reconstruct the high resolution image $\mathbf{z}$. First, the algorithm performs multiframe block matching to find the set of super-pixel tile shifts $\mathbf{R}_t(\underline{\mathbf{o}})$. The multiframe nature of the block matching reduces the matching artifacts common to pair-wise matching applied to low SNR data. Second, using the estimates $\hat{\mathbf{R}}_t(\underline{\mathbf{o}})$, we jointly register and reconstruct a high resolution blurry image by estimating the set of subpixel shifts $\underline{\mathbf{m}}_t$ and the image tile $\mathbf{z}_t$. This process is repeated for every tile to fully reconstruct the high resolution image $\mathbf{z}$.

A. Fast Multiscale Multiframe Block Matching

The first stage of the super-resolution algorithm involves multiframe block matching to estimate the set of unknown super-pixels shifts. The goal of this first processing stage is to produce a set of captured image tiles which contain only subpixel shifts. We achieve this by minimizing a cost function

of the form

$$J_1(\mathbf{p}_t, \underline{o}_t) = \sum_k \|\mathbf{R}_t^T(\mathbf{o}_k)\mathbf{y}_k - \mathbf{p}_t\|^2. \tag{2}$$

In this formulation, the vector $\mathbf{p}_t$ represents the unknown alias-free "reference" downsampled version of the image tile $\mathbf{z}_t$. We define $\mathbf{p}_t$ as

$$\mathbf{p}_t = \mathbf{DHz}_t, \tag{3}$$

where $\mathbf{H}$ represents an idealized low-pass filter to remove any aliasing artifacts associated with the downsampling operation. This multiframe estimation formulation is a joint estimation of both the alias-free downsampled tile and the super-pixel registration parameter. Since the first frame $\mathbf{y}_0$ defines the coordinate system, we are essentially looking for blocks in the set of images $\{\mathbf{y}_k\}$ which best fit an estimate of the tile $\mathbf{p}_t$. Since the captured images $\mathbf{y}_t$ are not free of aliasing artifacts, we implicitly treat the aliasing artifacts as noise. This assumption holds in practice because the strength of the alias-free signal is much stronger than that of the aliased content in digital mammogram images.

The cost functions of Eq. 2 falls into the class of estimation functionals known as separable nonlinear least squares (NLS) problems [9]. Such problems may be expressed as nonlinear search problems in a reduced space. For example, if we were to know the super-pixel motions $\underline{o}_t$ exactly, estimation of $\mathbf{p}_t$ would be reduced to a simple linear least squares (LS) problem where the estimate would be given by

$$\hat{\mathbf{p}}_t(\underline{o}_t) = \frac{1}{K+1}\left(\sum_k \mathbf{R}_t^T(\mathbf{o}_k)\mathbf{y}_k\right). \tag{4}$$

The vector $\mathbf{R}_t^T(\mathbf{o}_k)\mathbf{y}_k$ represents the cropped portion of the image $\mathbf{y}_k$ matching the unknown image tile region $\mathbf{p}_t$. Plugging this estimate back in to the original cost function, we obtain a cost function of only the unknown super-pixel motions

$$J_1(\underline{o}_t) = \sum_k \|\mathbf{R}_t^T(\mathbf{o}_k)\mathbf{y}_k - \hat{\mathbf{p}}_t(\mathbf{R}_t(\underline{o}))\|^2. \tag{5}$$

In this way, we have reduced the problem of joint registration and reconstruction to that of a registration problem with implicit reconstruction.

In practice, we impose a limit on the space of the pixel shifts based on geometric consideration of the imaging system. We assume the space of block shifts is constrained to a $[-W, W]$ search window around the reference frame image block $\mathbf{R}_t^T(0)\mathbf{y}_0$. This reduced search space significantly reduces the computational complexity of the multiframe block matching algorithm.

It is interesting to note that the traditional approach of estimating motions in a pairwise fashion is a simple approximation of this function where the estimate of the reference tile is given by $\hat{\mathbf{p}}_t = \mathbf{R}_t^T(0)\mathbf{y}_0$. Inserting this estimate into Eq. 2 produces a sum of K independent sub-optimization problems corresponding to the traditional pair-wise block matching approach. To give an idea of the performance degradation, we perform a simple Monte Carlo experiment using three one-dimensional signals. The three signals are all low-pass versions of a random signal having different low-pass filter cutoff frequencies controlling the amount of texture in the signal. The graph of Fig. 2 compares the registration error
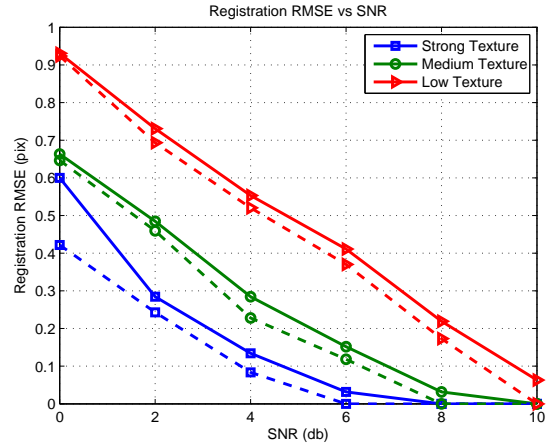

Fig. 2. The graph compares the average registration error (RMSE pixels) using three frames (two unknown super-pixel shifts) versus SNR, where the solid and dashed lines represent the pairwise and multiframe errors, respectively. We experimented on three random signals having increasing amounts of signal texture.

in estimating $\underline{o}$ by way of multiframe (5) or the pair-wise block matching. The simulation uses 500 samples for each data point. We use $K + 1 = 3$ frames for three different random signals with varied amount of signal texture. The performance difference demonstrates the loss incurred while estimating in a pairwise fashion as opposed to the optimal multiframe approach. This performance gap increases in case we use more frames. As we shall see experimentally, the inferior performance of pairwise approximation is significant for the digital mammography images.

B. Iterative Multiframe Sub-pixel Registration

After finding estimates of the super-pixel shifts $\hat{\mathbf{R}}_t^T(\underline{o})$, we address the problem of multiframe sub-pixel registration and reconstruction algorithm. We begin by using the super-pixel blocks estimates to crop the input images

$$\tilde{\mathbf{y}}_{k,t} = \mathbf{R}_t^T(\mathbf{o}_k)\mathbf{y}_k. \tag{6}$$

At this point, we (ideally) have a set of image tiles $\{\tilde{\mathbf{y}}_{k,t}\}$ containing only sub-pixel shifts related to the high resolution image tile according to

$$\tilde{\mathbf{y}}_{k,t} = \mathbf{DS}(\mathbf{m}_{k,t})\mathbf{z}_t + \mathbf{e}_t \tag{7}$$

The maximum likelihood (ML) estimate of both the set of sub-pixel shifts $\underline{m}_t$ and the unknown high resolution image tile $\mathbf{z}_t$ is computed as the minimum of

$$J_{2_{ML}}(\mathbf{z}_t, \underline{m}_t) = \sum_k \|\tilde{\mathbf{y}}_{k,t} - \mathbf{DS}(\mathbf{m}_{k,t})\mathbf{z}_t\|^2. \tag{8}$$

If we have some prior information about the unknown high resolution image $\mathbf{z}$, we may incorporate this into the cost

function to regularize the estimation problem. For instance, the unknown signal may be a Gaussian random variable with a covariance matrix $\mathbf{C_z}$. Assuming for the time being that the signal is zero-mean, the Maximum *a-posteriori* (MAP) estimate of the high resolution image and motion parameters is the minimum of the cost function

$$J_{2_{MAP}}(\mathbf{z}_t, \underline{m}_t) = \sum_k \|\tilde{\mathbf{y}}_{k,t} - \mathbf{DS}(\mathbf{m}_{k,t})\mathbf{z}_t\|^2 + \frac{\sigma^2}{\lambda}\mathbf{z}_t^T\mathbf{C}_{\mathbf{z}_t}^{-1}\mathbf{z}_t. \tag{9}$$

Since the cost functions of Eq. 8 and Eq. 9 are again separable NLS problems, we can perform a similar reduction as that in Sect. II-A. We focus on the MAP estimate henceforth where the estimate of the unknown image tile is

$$\hat{\mathbf{z}}_t(\underline{m} + t) = \mathbf{Q}^{-1}(\underline{m}_t)\mathbf{g}(\underline{m}_t) \tag{10}$$

where

$$\begin{aligned}
\mathbf{Q}(\underline{m}_t) &= \sum_k \mathbf{S}^T(\mathbf{m}_{k,t})\mathbf{D}^T\mathbf{DS}(\mathbf{m}_{k,t}) + \frac{\sigma^2}{\lambda}\mathbf{C}_{\mathbf{z}_t}^{-1} \\
\mathbf{g}(\underline{m}_t) &= \sum_k \mathbf{S}^T(\mathbf{m}_{k,t})\mathbf{D}^T\tilde{\mathbf{y}}_{k,t}
\end{aligned}$$

Plugging this form back into Eq. 8 we obtain a variable projections formulation [10]

$$J_2(\underline{m}_t) = \mathbf{g}^T(\underline{m}_t)\mathbf{Q}^{-1}(\underline{m}_t)\mathbf{g}(\underline{m}_t) \tag{11}$$

which must be maximized over the space of $\underline{m}_t$. At this point, we apply the numerical tricks originally described in [11] to speed up the optimization of the nonlinear cost function. We refer the reader to [11] for more information about this process. After running our optimization routine over the simplified cost function, we use Eq. 10 to obtain an estimate of the high resolution image tile.

III. EXPERIMENT

Our experimental imaging system is based on a modified Mammomat NovationDR digital mammography system. The system uses a stationary selenium-based detector having 85 μm pixels. Pixels with this size correspond to a Nyquist sampling rate of 5.6 line pairs per millimeter (lp/mm). We use a CIRS Model 011A breast phantom (CIRS Inc., Norfolk VA) to test out imaging algorithm. We introduce shifts in the image by two methods. First, we allow the X-ray tube to rotate by ± 1 degree. Second, we manually move the breast phantom to introduce motion into the system. The manual motion in the range of a few mm is completely uncontrolled. We acquire 15 frames at the low dosage level of 11.3 mAs at 28 kVp tube voltage. As a point of reference, we also acquire a single frame at the normal dosage of 226 mAs at 28 kVp tube voltage. We focus on the results of the test resolution chart to explore the contrast performance of the multiframe imaging system.

We apply our algorithm to 100×100 pixel tiles in the captured image to estimate 400×400 pixel high resolution images (enhancement $d = 4$). Figure 3 compares the performance of the pair-wise matching and sub-pixel registration with the multiframe matching and registration approaches.

The top-left image shows the reconstructed (factor $d = 4$) image where the both the super-pixel shifts $\underline{o}_t$ and the sub-pixel shifts $\underline{m}_t$ were estimated in a pairwise fashion. The image contains significant artifacts due to the errors in the super-pixel registration estimates. These errors occur because the pair-wise block matching is confused due to the periodic nature of the resolution chart and the low SNR. The top-right image shows the reconstructed image where the super-pixel shifts $\underline{o}_t$ were estimated using the multiframe approach and the sub-pixel shifts $\underline{m}_t$ were estimated in a pairwise fashion. The accurate multiframe block matching eliminates the coarse artifacts, but the image contains small sub-pixel motion error artifacts. The bottom-left image shows the reconstructed image when both the super-pixel and sub-pixel motions $\underline{o}_t$ and $\underline{m}_t$ were estimated using multiframe approaches. The images on the bottom-right show cropped portions of the resolution test chart. The multiframe sub-pixel registration algorithm better aligns the resolution bars.
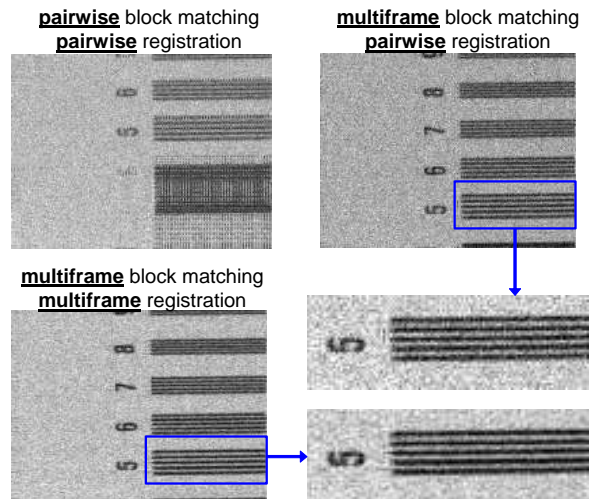

pairwise block matching **pairwise** registration   **multiframe** block matching **pairwise** registration

multiframe block matching **multiframe** registration

Fig. 3. The top-left image shows the reconstructed (factor $d = 4$) image where the both the super-pixel shifts $\underline{o}_t$ and the sub-pixel shifts $\underline{m}_t$ were estimated in a pairwise fashion, containing artifacts due to the errors in the super-pixel registration estimates. The top-right image shows the reconstructed image where the super-pixel shifts $\underline{o}_t$ were estimated using the multiframe approach and the sub-pixel shifts $\underline{m}_t$ were estimated in a pairwise fashion. The accurate multiframe block matching eliminates the coarse artifacts, but the image contains small sub-pixel motion error artifacts. The bottom-left image shows the reconstructed image when both the super-pixel and sub-pixel motions $\underline{o}_t$ and $\underline{m}_t$ were estimated using multiframe approaches. The images on the bottom-right show cropped portions of the resolution test chart. The multiframe sub-pixel registration algorithm better aligns the resolution bars.

Figure 4 gives a visual example of the performance of the multiframe reconstruction algorithm in terms of resolution enhancement. The image on the top left (a) shows a single low resolution image tile for an input low dosage (11.3 mAs) image. The image on the top right (b) shows an image taken at normal dosage (226 mAs) linearly interpolated by a factor of $4\times$. At the improved SNR, we observe the aliasing artifacts present at the 9 and 10 lp/mm resolution bars. The bottom left image (c) shows the resulting estimate $\hat{\mathbf{p}}_t$ after motion

compensating and averaging the 15 frames. Interestingly, we observe that the image appears to have better SNR than the single normal dosage image even though the cumulative dosage of the multiframe image is only 169.5 mAs. The image appears slightly blurry, however, due to averaging over the sub-pixel shifts. This suggests that the multiframe imaging might also provide an interesting contrast vs SNR tradeoff variable as opposed to the standard dosage control. The bottom left image (d) shows the estimate of the high resolution tile $\hat{\mathbf{z}}_t$ after applying the full multiframe registration and reconstruction algorithm. The aliasing artifacts are eliminated and contrast is restored above the Nyquist rate.
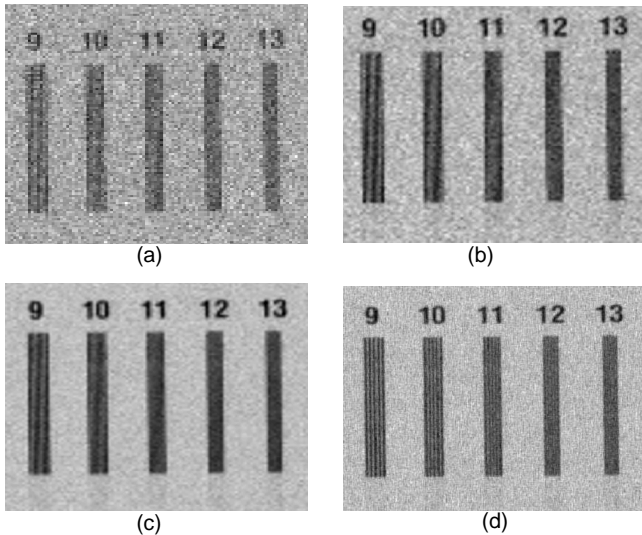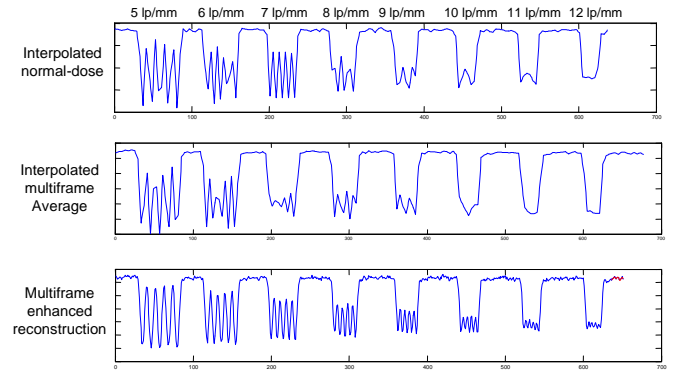

Fig. 5. The three curves show slices through the horizontal resolution bar charts. The sampling rate of the system corresponds to around 5.6 lp/mm. The top curve shows a slice through the normal dosage image interpolated by a factor of four. We clearly observe evidence of aliasing for the spatial frequencies above this sample rate. This becomes more apparent as we look at the middle slice through the interpolated motion compensated average. In this slice the contrast above the sampling rate is destroyed by averaging out the aliasing artifacts. The bottom slice shows the slice through the multiframe reconstruction. The slice clearly shows contrast preservation beyond $2\times$ the sample rate of the system.

Fig. 4. (a) low-resolution low-dosage image $\mathbf{y}_{0,t}$ (3 db). (b) Interpolated normal dosage image (12.4 db). (c) Interpolated average of 15 super-pixel registered frames $\hat{\mathbf{p}}$ (13 db). (d) Multiframe reconstructed image $\hat{\mathbf{z}}_t$ (11.2 db).

To get an idea of the effective modulation transfer function (MTF) of the multiframe system, we plot slices through the full resolution test chart region in Fig. 5. The top curve shows the slice through an interpolated normal dosage image tile. The middle curve shows the slice through the multiframe average image $\hat{\mathbf{p}}$. In both of these cases, we observe the apparent aliasing beyond the Nyquist sampling frequency of 5.6 lp/mm. The bottom graph shows a slice through the multiframe resolution enhanced image $\hat{\mathbf{z}}$. The system exhibits alias-free sampling out to more than $2\times$ the native sampling limit of the detector. The contrast drops low enough to fall into the noise around 11 or 12 lp/mm

IV. CONCLUSION

In this paper, we have proposed a novel method for acquiring high resolution X-ray images by combining multiple lower resolution images. The proposed multiframe registration algorithm demonstrates superior performance over standard pair-wise registration techniques in this low SNR regime. The experimental results confirm that multiframe imaging can provide a new tradeoff in the SNR versus resolution tradeoff for digital mammography. In the future, we explore the fundamental tradeoffs between radiation exposure, number of frames, and reconstruction performance.

REFERENCES

[1] E. Pisano, C. Gatsonis, E. Hendrick, M. Yaffe, J. Baum, S. Acharyya, E. Conant, L. Fajardo, L. Bassett, C. D'Orsi, R. Jong, and M. Rebner, "Diagnostic performance of digital versus film mammography for breast cancer screening- the results of the american college of radiology imaging network (ACRIN) digital mammographic imaging screening trial (DMIST)," *New England Journal of Medicine*, October 2005.

[2] S. Farsiu, D. Robinson, M. Elad, and P. Milanfar, "Advances and challenges in super-resolution," *International Journal of Imaging Systems and Technology*, vol. 14, no. 2, pp. 47–57, October 2004.

[3] S. Park, M. Park, and M. G. Kang, "Super-resolution image reconstruction, a technical overview," *IEEE Signal Processing Magazine*, vol. 20, no. 3, pp. 21–36, May 2003.

[4] P. Bernhardt, M. Lend, and F. Deinzer, "New technologies to reduce pediatric radiation doses." *Pediatric Radiology*, vol. 36, pp. 212–215, Sept. 2006.

[5] D. Robinson and P. Milanfar, "Statistical performance analysis of super-resolution," *IEEE Transactions on Image Processing*, vol. 15, June 2006.

[6] S. Farsiu, D. Robinson, M. Elad, and P. Milanfar, "Fast and robust multi-frame super-resolution," *IEEE Transactions on Image Processing*, vol. 13, no. 10, pp. 1327–1344, October 2004.

[7] S. Farsiu, M. Elad, and P. Milanfar, "Constrained, globally optimal, multi-frame motion estimation," in *Proc. of the 2005 IEEE Workshop on Statistical Signal Processing*, July 2005, pp. 1396–1401.

[8] P. Vandewalle, S. Susstrunk, and M. Vetterli, "A frequency domain approach to registeration of aliased images with application to super-resolution," *EURASIP Journal on Applied Signal Processing*, 2006.

[9] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice Hall Inc., 1993.

[10] G. Golub and V. Pereyra, "Separable nonlinear least squares: the variable projection method and its applications," *Institute of Physics Inverse Problems*, vol. 19, pp. R1–R26, 2003.

[11] M. D. Robinson, S. Farsiu, and P. Milanfar, "Optimal registration of aliased images using variable projection with applications to super-resolution," *The Computer Journal*, Advance online publication, doi: 10.1093/comjnl/bxm007, 2007.