

Robinson M.D, Chiu S.J., Lo J.Y., Toth C.A., Izatt J.A, Farsiu S,
"Novel Applications of Super-Resolution in Medical Imaging",
Book Chapter in *Super-Resolution Imaging*, Peyman Milanfar
(Editor). CRC Press, Pages 383-412, 2010

1

New Applications of Super-resolution in Medical Imaging

M. Dirk Robinson

Ricoh Innovations Inc., Menlo Park, CA, USA

Stephanie J. Chiu, Cynthia A. Toth

Duke University, Durham, NC, USA

Joseph A. Izatt, Joseph Y. Lo

Duke University, Durham, NC, USA

Sina Farsiu

Duke University, Durham, NC, USA

CONTENTS

1.1	Introduction	2
1.2	The Super-Resolution Framework	3
1.2.1	Image Capture Model	3
1.2.2	Super-Resolution Estimation Framework	4
1.3	New Medical Imaging Applications	6
1.3.1	Super-Resolution in Low Radiation Digital X-ray Mammography	6
1.3.1.1	Multiframe Shift Estimation	7
1.3.1.2	Multiframe ForWarD Deconvolution and Denoising	10
1.3.1.3	Experimental X-ray Results	12
1.3.2	Super-Resolution in Optical Coherence Tomography	15
1.3.2.1	Proposed Method: Sparse Repeated Imaging	16
1.3.2.2	Multiframe Joint Registration	19
1.3.2.3	Experimental Results	21
1.4	Conclusion	22
1.5	Acknowledgment	23

The image processing algorithms collectively known as *super-resolution* have proven effective in producing high-quality imagery from a collection of low-resolution photographic images. In this chapter, we examine some of the advantages and challenges of applying the super-resolution framework to applications in medical imaging. We describe two novel applications in detail. The first application addresses the problem of improving the quality of digital mammography imaging systems while reducing X-ray radiation exposure. The

second application addresses the problem of improving the spatio-temporal resolution of spectral domain optical coherence tomography systems in the presence of uncontrollable patient motion. Experimental results on real data sets confirm the effectiveness of the proposed methodologies.

1.1 Introduction

The invention of the charge coupled device (CCD) created a new era of imaging wherein optical images could be efficiently captured by an array of solid-state detectors and stored as digital information. The resolution of the captured image depended on the size and number of these detectors. Most imaging applications critically depend on high-resolution imagery. Increasing resolution by improving detector array resolution is not always a feasible approach to improving resolution. For example, while improvements in semiconductor manufacturing have translated into higher-resolution image sensors, shrinking pixel sizes has a tendency to decrease signal-to-noise ratios (SNR) and light sensitivity. Furthermore, practical cost and physical limitations limit the ability to change detectors for most legacy imaging systems. To address this issue, the image processing community is developing a collection of algorithms known as *super-resolution* for generating high-resolution imagery from systems having lower-resolution imaging detectors. These algorithms combine a collection of low-resolution images containing aliasing artifacts and restore a high-resolution image. The ability to transcend the fundamental resolution limits of sensors using super-resolution algorithms has shown significant progress and capability in the area of photographic imaging. By far, the majority of applications using super-resolution technology have been in the area of photographic imagery for either consumer or defense-type applications, which are discussed in the other chapters of this book.

Relatively recently, researchers have begun developing methods to extend the super-resolution framework to different medical imaging applications. Medical imaging applications differ from photographic imaging in several key respects. First, unlike photographic imaging, medical imaging applications often use highly controlled illumination of the human subject during image acquisition. As with any imaging system, stronger illumination energy results in higher signal-to-noise ratios. In the case of medical imaging, however, illumination radiation is limited to prevent tissue damage, thereby limiting the SNR to well below that of photographic imaging. Second, imaging speed is more important in medical imaging applications than in photographic applications. Short acquisition times both limit patient discomfort and minimize imaging artifacts associated with patient movement. Third, unlike photographic imaging, the goal of medical imaging is to facilitate the detection or diagnosis of disease, rather than produce visually pleasing imagery. Consequently, im-

age processing artifacts are much less tolerable in medical images than in photographic applications. Luckily, medical imaging systems operate under highly controlled environments with highly similar objects. Algorithm developers can leverage prior knowledge about the anatomy or biology to improve image quality. Finally, the majority of medical imaging applications involve creating images from radiation propagation through three-dimensional objects. Thus, while the final images are two-dimensional, they represent some form of projection through a three-dimensional volume.

In this chapter, we describe super-resolution and its applications from the medical imaging community's point of view. In Section 1.2, we describe the general super-resolution framework and provide a brief review of the different super-resolution algorithms. In Section 1.3.1, we introduce the first of two novel applications of super-resolution in medical imaging. Namely, we describe how we tailor the super-resolution framework to improve the resolution for digital X-ray mammography. In Section 1.3.2, we describe how we apply the super-resolution framework to Optical Coherence Tomography (OCT). Finally, we conclude in Section 1.4 with some thoughts about future applications of super-resolution in medical imaging.

1.2 The Super-Resolution Framework

The goal of super-resolution image processing is to extract a high-resolution image from a collection of images containing aliasing artifacts. When a collection of aliased, low-resolution images contains sufficient variation, the high-resolution, aliased image content can be separated from the low-resolution image content thereby increasing the image resolution. This type of super-resolution is not to be confused with optical methods for transcending the optical diffraction limit (e.g. [63]). There are a number of broad reviews of super-resolution algorithms [3, 15, 35]. In this section, we describe the general super-resolution imaging framework. The section begins with a description of a generic image capture model and concludes with a general super-resolution estimation framework.

1.2.1 Image Capture Model

The image capture model describes the various physical processes involved when capturing a set of images. As with most multiframe image super-resolution algorithms, the collection of images must contain relative motion between the sets of images from which resolution is enhanced. We assume very simple translational motion for several reasons. First, even if this motion is not appropriate in a wide-field sense, the motion model is typically accurate for local regions within the images [1]. Second, the imaging acquisition system

can often be controlled to induce only translational motion in the captured images. Thus, for the remainder of this chapter, bear in mind that when we refer to an image, we could also be referring to a cropped portion of a larger image. This model has been used in several previous works on super-resolution [11, 16, 23, 32].

The general image capture model, or forward model, combines the various effects of the digital image acquisition process such as point-wise blurring, motion, undersampling, and measurement noise. We represent the forward imaging model using matrix notation as

$$\mathbf{y}_k = \mathbf{DHS}(\mathbf{v}_k)\mathbf{x} + \mathbf{e}_k, \quad (1.1)$$

where \mathbf{x} and \mathbf{y} are rearranged in lexicographic order. Here, the vector \mathbf{y}_k represents $B \times B$ (assumed square without loss of generality) samples of the captured image $y_k(m'_1, m'_2)$, where $m'_1, m'_2 \in [0, B - 1]$, are ordered as a $(B)^2 \times 1$ vector. The captured image is undersampled with respect to an unknown high-resolution image $x(m_1, m_2)$, where $m_i \in [0, D_i B - 1]$, by a factor of D_1 and D_2 in each of the two respective dimensions. The vector \mathbf{x} represents samples of the unknown $D_1 B \times D_2 B$ high-resolution image tile $x(m_1, m_2)$ similarly ordered. The warping operator $\mathbf{S}(\mathbf{v}_k)$ of size $D_1 D_2 B^2 \times D_1 D_2 B^2$ represents the subpixel spatial shifts between the captured images. Without loss of generality, we assume that the 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 the remaining K images. Note that, here to simplify notations, instead of $\mathbf{v}_{k,0}$, we use $\mathbf{v}_k = [v_{k_1}, v_{k_2}]$, which is the spatial shifting between the reference frame (0th) and the k th frame. In Section 1.3.2, however, we will use the full form of $\mathbf{v}_{k,i}$ to represent the motion between the i th and k th frames. In our model, we assume that these spatial shifts are continuous values in the range of $[-D_i, D_i]$. This corresponds to the range of subpixel motion in the captured images. The downsampling operator \mathbf{D} of size $B^2 \times D_1 D_2 B^2$ captures the undersampling of the detector. The matrix \mathbf{H} represents the blurring associated with the imaging system. This blurring could be the result of multiple processes within the imaging system. For example, this blurring could be the result of integration apertures or motion during the image capture, or scattering of radiation in the object medium as a point spread function (PSF). For the time being, we will assume that this can be reliably measured or estimated from some other process (note [39] as an example of jointly estimating the high-resolution image and the blur parameters in a Bayesian framework). Finally, \mathbf{e}_k of size $B^2 \times 1$ 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 . This model is sufficiently broad as to cover a wide variety of imaging systems.

1.2.2 Super-Resolution Estimation Framework

The goal of super-resolution image processing is to estimate the high-resolution image \mathbf{x} from the set of captured images $\{\mathbf{y}_k\}$. The most common estimation framework begins with a cost function or penalty function relating the observed data to the unknown high-resolution image. The most common statistical framework found in super-resolution is that of the *maximum a posteriori* (MAP) penalty function of the form

$$\Omega(\mathbf{x}, \{\mathbf{v}_k\}) = \Omega_d(\mathbf{x}, \{\mathbf{v}_k\}) + \Omega_p(\mathbf{x}). \quad (1.2)$$

The MAP functionals are based on the construction of a cost function (Ω), which is the summation of two distinct terms. One is the data penalty term Ω_d , which measures the closeness of data to the estimates. The other is the regularization term Ω_p , which applies the prior information about or constraints on the unknown high-resolution image (\mathbf{x}).

Early MAP functionals used in super-resolution processing utilized simple quadratic data penalty and regularization terms [10, 45]. The most commonly employed regularization terms use Tikhonov type functionals despite their tendencies to reduce edge contrast. These quadratic regularization functionals penalize the amount of high spatial-frequency energy in the high-resolution image estimate. For example, using the generic imaging model of Equation (1.1), a typical quadratic MAP functional takes the form

$$\Omega(\mathbf{x}, \{\mathbf{v}_k\}) = \sum_{k=0}^K \|\mathbf{y}_k - \mathbf{DHS}(\mathbf{v}_k)\mathbf{x}\|^2 + \lambda \mathbf{x}^T C_{\mathbf{x}}^{-1} \mathbf{x}, \quad (1.3)$$

where $C_{\mathbf{x}}^{-1}$ is often a spatial high-pass operator and λ is the weighting scalar. When $C_{\mathbf{x}}$ is the exact covariance of the unknown high-resolution image, then this cost function produces the ideal Wiener filter estimate of the unknown image. This MAP functional has the advantage of being quadratic, which means that the penalty function has an analytic solution that is a linear function of the input measurements.

Through the years, application of more advanced prior functions Ω_p such as Adaptive Kernel regression [53] and Bilateral Total-Variation (B-TV) [16] have produced higher quality estimates of the final by imposing more accurate assumptions about the image content. The tradeoff, however, is that such non-linear prior functionals are more expensive to evaluate and require more computationally-complex iterative minimization techniques. For example, the B-TV cost function is defined as

$$\Omega_{\text{B-TV}}(\mathbf{x}, \{\mathbf{v}_k\}) = \sum_{k=0}^K \|\mathbf{y}_k - \mathbf{DHS}(\mathbf{v}_k)\mathbf{x}\|_2^2 + \lambda \sum_{t_1, t_2=-L}^L \varrho^{|t_1|+|t_2|} \|\mathbf{x} - \mathbf{S}(\mathbf{t})\mathbf{x}\|_1, \quad (1.4)$$

where $\mathbf{t} = [t_1, t_2]$ is a set of integer pixel shifts and $0 < \varrho \leq 1$ is a constant

[16]. Such non-quadratic functionals can, however, preserve many important features of images such as edges. Also, MAP-based robust super-resolution techniques (e.g. [9, 16, 36]) are able to reduce the effect of outliers such as motion estimation error.

Both the quadratic and non-quadratic MAP functionals require knowledge of the relative shifts between the collection of low-resolution images. When the SNR is reasonably high and the amount of aliasing artifacts are low, then the shifting parameters $\{\mathbf{v}_k\}$ can be reasonably estimated in an initial step from the captured images \mathbf{y}_k . Theory as well as experimental evidence, however, suggests that using a separate shift estimation process in low-SNR cases is suboptimal. Therefore, the critical issue of joint super-resolution and motion estimation problem has been the topic of several papers (e.g. [2, 24, 43, 58, 59, 62]). Note that, additional priors on motion vector distribution may be also added to the above cost function [24].

1.3 New Medical Imaging Applications

Early, fast, and accurate detection of imaging biomarkers of the onset and progression of diseases is of great importance to the medical community since early detection and intervention often results in optimal treatment and recovery. The advent of novel imaging systems has for the first time enabled clinicians and medical researchers to visualize the anatomical substructures, pathology, and functional features *in vivo*. However, earlier biomarkers of disease onset are often critically smaller or weaker in contrast compared to their corresponding features in the advanced stages of disease. Therefore, medical imaging community strives for inventing higher-resolution/contrast imaging systems. As noted in Section 1.2, super-resolution can be beneficial in improving the image quality of many medical imaging systems without the need for significant hardware alternation.

An excellent review of previous medical imaging applications of super-resolution is given in [20]. We refer the interested reader to [20] for a broad review of applications in magnetic resonance imaging (MRI) [21, 38], functional MRI (fMRI)[37], and positron emission tomography imaging system (PET)[29, 30]. In the following two sections, we explore novel applications of the super-resolution framework to medical imaging. The first application is in the area of X-ray digital mammography. The second is in the area of Optical Coherence Tomography (OCT). Each application has its own unique properties that demand customization of the general super-resolution framework described in the previous section. In both applications, the advantage of applying the super-resolution framework is achieved by special modification of the standard image acquisition technique.

1.3.1 Super-Resolution in Low Radiation Digital X-ray Mammography

Digital mammography provides the opportunity to efficiently control and capture digital images of the breast, while exposing the patient to the minimum amount of radiation. Today's digital detectors cannot shrink pixel sizes to increase resolution without sacrificing the SNR measurement. To maximize image resolution, we have explored digitally combining multiple low-dosage images, each containing spatial shifts. These shifts are the result of patient movement, intentional dithering of the detector, vibration in the imaging system, and small movement of the imaging gantry. In practice, the motion contained in the captured images is a combination of all such sources necessitating accurate registration of the aliased, low-resolution images.


Applying super-resolution processing to X-ray imaging requires overcoming two challenges. The first is the large amount of data associated with digital mammogram images. The captured low-resolution images could have as much as 10 megapixels worth of data. Thus, computational efficiency is extremely important during processing. Second, the total radiation exposure over the collection of images cannot exceed that of a normal X-ray image dosage. Therefore, the captured data has extremely low peak-SNR (PSNR). For example, Figure 1.1 compares a high-dosage X-ray image (computed PSNR¹ \simeq 13 dB) with the very low exposure images (computed SNR \simeq 3 dB) used in our multiframe scheme.

Thus, providing high-resolution imagery requires sophisticated, nonlinear reconstruction techniques to address the extremely low SNR of the captured images. To address these two challenges, we apply a divide and conquer approach to both improve efficiency while maximizing the denoising effectiveness. To achieve this, we propose a three-stage (registration, reconstruction, and restoration) algorithm. The overall algorithm procedure is shown in Figure 1.2. The entire algorithm operates on a tile-based fashion. The process begins by finding a collection of tiles with approximately equal regions-of-interest. Then, each of these tiles are registered to a subpixel precision to estimate the shifts $\{\mathbf{v}_k\}$. Next, we apply a multiframe image restoration step with a weak quadratic prior function, resulting in a deblurred aliasing free image with reconstruction artifacts with known statistics. Next, we perform a fast estimate of wavelet coefficients, which best match the reconstruction artifacts in the previous step. Finally, we apply a nonlinear wavelet thresholding-based denoising step, to recover an efficiently denoised super-resolved image. In what follows we describe each step in detail.


¹In this work, the PSNR was computed numerically as $\text{PSNR} = 20 \log_{10} \frac{s}{n}$. In experiments on real images s is the grayscale difference between the minimum and maximum signal regions and n is the noise standard deviation estimated from flat regions. In simulated experiments, n is the RMSE error between the estimated and ground truth image.



(a)



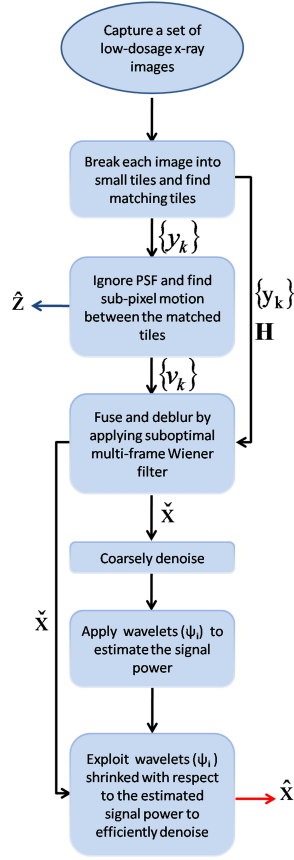
(b)



(c)

FIGURE 1.1

Mammogram X-ray images from the phantom breast in (a). The red rectangular section in (a) is zoomed in (b) and (c). The high dosage image in (b) is captured at 226mAs ($\text{PSNR} \approx 13\text{dB}$). The extremely low-dosage image in (c) is captured at 11.3 milliAmpere-second (mAs) ($\text{PSNR} \approx 3\text{dB}$). Regardless of SNR, both images show aliasing artifacts.

**FIGURE 1.2**

The block diagram shows the non-iterative super-resolution algorithm we apply to digital mammogram images.

1.3.1.1 Multiframe Shift Estimation

Apart from the basic block-matching required to find a collection of approximately registered image tiles $\{y_k\}$ [41], the super-resolution algorithm begins with a multiframe subpixel shift estimation algorithm. The efficiency of this stage is improved by ignoring the optical blur. Considering the locally space-invariant PSF and shift assumptions in our models, we may reverse the order of the shifting and blur operators in Equation (1.1) [11] and reformulate the forward model without the blur operator as

$$y_k = \mathbf{DS}(v_k)z + e_k, \quad (1.5)$$

where $z = \mathbf{H}x$ is the unknown high-resolution blurry image.

The registration algorithm begins with a variant of the quadratic penalty function of Equation (1.3). The optimization process then will be formulated as

$$\Omega_1(\mathbf{z}, \{\mathbf{v}_k\}) = \sum_{k=0}^K \|\mathbf{y}_k - \mathbf{D}\mathbf{S}(\mathbf{v}_k)\mathbf{z}\|_2^2 + \lambda \mathbf{z}^T \mathbf{C}_z^{-1} \mathbf{z}, \quad (1.6)$$

where \mathbf{C}_z is the covariance matrix of the unknown signal \mathbf{z} , which is typically assumed to be stationary. In other words, we can ignore the optical blur for this stage of processing. A typical solution to the above problem is the cyclic coordinate-descent method [24], in which in each iteration one unknown variable is updated based on the estimate of the other unknown variable from the previous iteration.

Noting that Equation (1.6) is known in numerical analysis literature as the Separable Non-linear Least Squares problem [18], we momentarily assume in our Variable-Projection technique [42, 58] that the non-linear parameters (motion-vectors) are known. Consequently, the estimate of the set of linear parameters (\mathbf{z}) is computed as

$$\hat{\mathbf{z}} = (\mathbf{Q}(\{\mathbf{v}_k\}) + \lambda \mathbf{C}_z^{-1})^{-1} \mathbf{g}(\{\mathbf{v}_k\}), \quad (1.7)$$

where

$$\mathbf{Q}(\{\mathbf{v}_k\}) = \frac{1}{\sigma^2} \sum_{k=0}^K \mathbf{S}^T(\mathbf{v}_k) \mathbf{D}^T \mathbf{D} \mathbf{S}(\mathbf{v}_k), \quad (1.8)$$

$$\mathbf{g}(\{\mathbf{v}_k\}) = \frac{1}{\sigma^2} \sum_{k=0}^K \mathbf{S}^T(\mathbf{v}_k) \mathbf{D}^T \mathbf{y}_k. \quad (1.9)$$

We plug the parametric estimate of the blurry high-resolution image ($\hat{\mathbf{z}}$) into the MAP functional (Eq. (1.6)) and after some algebraic simplifications, we get a new (maximization) cost function that only relies on the motion-vectors:

$$\Omega_1(\{\mathbf{v}_k\}) = \mathbf{g}(\{\mathbf{v}_k\})^T (\mathbf{Q}(\{\mathbf{v}_k\}) + \lambda \mathbf{C}_z^{-1})^{-1} \mathbf{g}(\{\mathbf{v}_k\}). \quad (1.10)$$

Note that, unlike the cyclic coordinate-descent method, we require no iterations between the sets of parameters since we do not explicitly calculate Equation (1.7). Indeed, a direct approach to maximize Equation (1.10) involves inverting a large matrix of size $D_1 D_2 B^2 \times D_1 D_2 B^2$ which is computationally challenging for even small image tiles. In [42], we described a series of numerical tricks to speed up the process. One is solving the problem in the Fourier domain and taking advantage of the spectral folding phenomenon in aliased images.

1.3.1.2 Multiframe ForWarD Deconvolution and Denoising

The output of the previous algorithm is an estimate of the set of sampling shift offsets $\{\mathbf{v}_k\}$ with which we can estimate the high-resolution image \mathbf{x} . We estimate a high-quality super-resolution image using a non-iterative, two-stage, linear deconvolution and nonlinear denoising algorithm. The algorithm addresses the SNR versus sharpness tradeoff inherent to quadratic-type regularization functionals without resorting to iterative, nonlinear regularization penalty functions. More information about the algorithm can be found in [44].

The first stage of the algorithm performs multiframe deconvolution using a weak quadratic penalty function. Armed with estimates of the image shifts, a sharpened, high-resolution image can be obtained using a variant of Equation (1.7) given by

$$\tilde{\mathbf{x}} = \mathbf{B}^{-1}(\hat{\mathbf{v}})\mathbf{H}^T\mathbf{g}(\hat{\mathbf{v}}), \quad (1.11)$$

where

$$\mathbf{B}(\hat{\mathbf{v}}) = \mathbf{H}^T\mathbf{Q}(\hat{\mathbf{v}})\mathbf{H} + \lambda\mathbf{C}_{\mathbf{x}}^{-1}, \quad (1.12)$$

$\hat{\mathbf{v}} = [\mathbf{v}_1, \dots, \mathbf{v}_k]^T$ and \mathbf{Q} and \mathbf{g} were defined in Equations (1.8) and (1.9).

In this first stage of the algorithm, we use a very small value of λ so as to under-regularize the estimate of the high-resolution image estimate $\hat{\mathbf{x}}$. This creates a very sharp high-resolution image at the expense of extreme noise amplification. The second stage of the algorithm involves eliminating the noise while preserving the image signal content. We achieve this with a type of wavelet thresholding algorithm similar to the ForWard algorithm of [34] or the BayesShrink algorithm of [4]. The wavelet thresholding algorithm operates by first applying a wavelet transform to the noisy high-resolution image represented as

$$\mathbf{w} = \Psi\tilde{\mathbf{x}}, \quad (1.13)$$

where the matrix Ψ represents the wavelet transform operator and \mathbf{w} the wavelet coefficients. Then, the wavelet coefficients are scaled according to

$$w'_i = \text{sgn}(w_i)\max(0, |w_i| - \gamma_i), \quad (1.14)$$

where w_i represents the individual wavelet coefficients, sgn is the sign function, and γ_i represents the threshold value for those wavelet coefficients. After applying this threshold, the inverse wavelet transform is applied to the thresholded wavelet coefficients to get the final denoised estimate of the high-resolution image

$$\hat{\mathbf{x}} = \Psi^{-1}\mathbf{w}'. \quad (1.15)$$

This type of wavelet thresholding has the ability to eliminate noise while preserving signal content. More information about how this is implemented efficiently can be found in [44].

1.3.1.3 Experimental X-ray Results

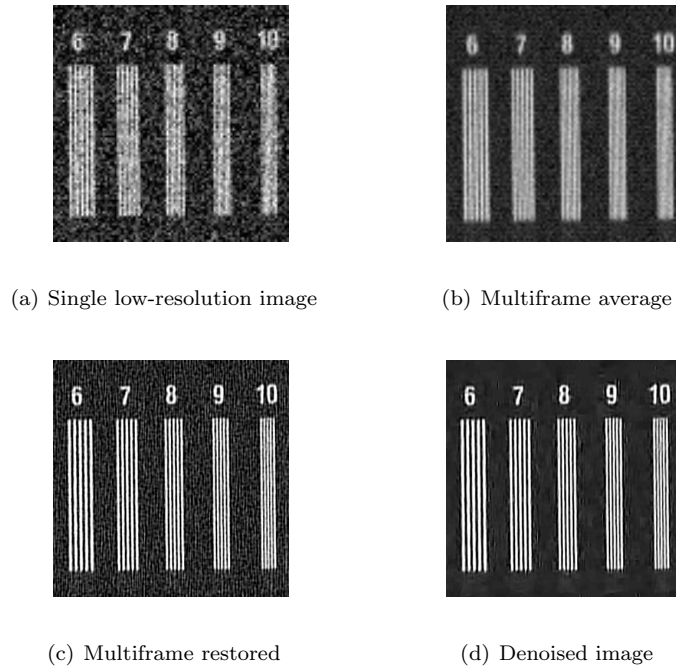
We applied our multiframe reconstruction and restoration algorithm to real images captured on an experimental X-ray imaging system. Our experimental imaging system is based on a Mammomat NovationTOMO digital mammography prototype system (Siemens Medical Solutions, Erlangen Germany)², stationed at Duke University Medical Center. The system uses a stationary selenium-based detector having $85 \mu\text{m}$ pixels. Pixels with this size correspond to a Nyquist sampling rate of 5.6 line pairs per millimeter (lp/mm). We used a CIRS model 11A breast phantom (CIRS Inc., Norfolk VA) to test our super-resolution algorithms. We introduced shifts in the image by two methods. First, we allowed the X-ray tube to rotate by ± 1 degree. Second, we manually moved the breast phantom to introduce motion into the system. This manual motion was completely uncontrolled. Our dataset consisted of 15 frames at the low dosage level of 11.3 mAs at 28 kVp tube voltage. As a point of reference, we also acquired a single frame at a more typical dosage of 226 mAs at 28 kVp tube voltage (Fig. 1.1), which is 25.

The breast phantom includes several testing features including a pair of resolution bar charts and small grains that mimic calcification in the breast. The results reported here are focused on the test resolution chart and the calcification grains that best represent the contrast performance and potential improved-detection abilities of the multiframe image reconstruction system.

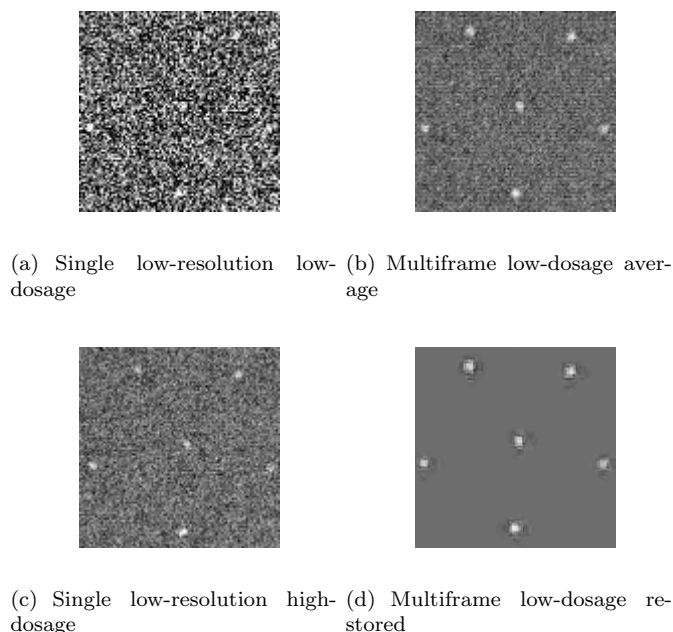
We applied our algorithm to 100×100 pixel tiles in the captured image to estimate 400×400 pixel high-resolution images (enhancement $D = 4$). We modeled our system PSF as a heavy-tailed exponential energy distribution with $\beta = 1.5$. To get a measure of the PSNR, we calculated the standard deviation in a textureless region of the phantom. We also measured the difference in grayscale values between registration bars in the resolution chart to get an approximate PSNR value of 3 dB. We employed 2-tap Daubechies filters for the soft thresholding wavelet functions and 6-tap Daubechies filters for the coarse denoising by way of hard wavelet coefficient thresholding. We focus on the portion of the resolution chart beyond the Nyquist rate for the imaging system (5.6 lp/mm). The numbers indicate the resolution in terms of line pairs per millimeter (lp/mm). Figure 1.3 shows the images throughout the super-resolution process.

The first image (Fig. 1.3(a)) shows one of the 15 low-dosage images. The image has very low SNR and shows some of the aliasing associated with an undersampled detector. The second image (Fig. 1.3(b)) shows an example of

²Caution: Investigational Device. Limited by US Federal law to investigational use. The information about this product is preliminary. The product is under development and is not commercially available in the US; and its future availability cannot be ensured.

**FIGURE 1.3**

Different restoration techniques applied on the low-dosage set of images. (a) Low-dosage image, (b) Multiframe averaged image, (c) Multiframe restored $\hat{\mathbf{x}}$, (d) Denoised super-resolved image $\hat{\mathbf{x}}$. The multiframe average image shows the aliasing present in the captured image. The super-resolved images show image contrast beyond the native sampling rate of the system. The total dosage of using 15 of these frames ($15 \times 11.3 = 170\text{mAs}$) is still less than the high dosage image of 225 mAs in Figure 1.1(b) with clear aliasing artifacts on resolution bars labeled with numbers higher than 8.

**FIGURE 1.4**

Different restoration techniques applied on the low-dosage set of images. (a) Low-dosage image, (b) Multiframe averaged image, (c) High-dosage image at 226mAs, (d) Denoised image \hat{x} ($15 \times 11.3 = 170\text{mAs}$). Restoration combining the 15 low-dosage frames in (d) frames, most clearly demonstrating the pentagram-shaped set of micro-calcification cluster.

averaging the 15 low-resolution frames followed by interpolation. While the SNR is improved, the aliasing contained in the low-resolution images becomes clear as the 5 bars appear as three bars above 7 lp/mm. The third image (Fig. 1.3(c)) shows the resulting image \hat{x} after applying the multiframe image restoration step. This image shows contrast improvement but at the expense of significant noise amplification. The final image (Fig. 1.3(d)) shows the final image estimate \hat{x} after applying the non-linear wavelet thresholding denoising algorithm. The image shows that the contrast is preserved while eliminating most of the amplified noise.

The primary goal of mammography is detecting and diagnosing cancerous lesions in the breast. The breast phantom includes small grains of calcium for evaluating the diagnostic capability of micro-calcifications. Figure 1.4 shows another block from the same experiment demonstrating the ability of the nonlinear denoising algorithm to clearly eliminate noise while preserving the

signal for a cluster of 0.275 mm sized calcite grains. This provides confidence in the nonlinear denoising algorithm's ability to discern signal from noise.

1.3.2 Super-Resolution in Optical Coherence Tomography

The invention and utilization of Optical Coherence Tomography (OCT), first reported in [27] in 1991, has had a profound impact on many fields of research, especially ophthalmic sciences [25, 26, 46, 49, 55]. OCT systems provide non-invasive yet high-resolution *in-vivo* images of retinal structures, which can be used to define imaging biomarkers of the onset and progression of many ophthalmic diseases. By employing an interferometer [8, 17], several OCT based imaging systems have been developed through out the years, most notably the time-domain OCT (TDOCT) and ultrahigh resolution OCT (UHROCT) [40]. The advent of the Spectral Domain Optical Coherence Tomography (SDOCT) system has further improved the image quality and acquisition speed [33, 51, 60, 61, 64, 65]. Today, several commercial SDOCT systems are available with similar capabilities and 20-30 kHz A-scan rates.

A noninvasive, accurate characterization of retinal lesions and other pathological abnormalities is only possible with high-resolution, 3-D ocular imaging. Commonly, the lateral, axial, and azimuthal resolution of many imaging systems, including OCT, are associated with (and calculated based on) the illumination source characteristics (e.g. bandwidth), the optical path (e.g. diffraction limit due to pupil diameter, ocular aberrations, dispersion, etc.), and other physical characteristics. Utilization of fast and efficient CCD detectors has facilitated the creation of aliasing-free 3-D images of anatomical structures. However, for some *in vivo* imaging applications, the SDOCT acquisition time is not short enough to avoid abrupt motions such as blinking, thus creating motion artifacts in the densely sampled volumetric measurements (Fig. 1.6). Therefore, in practice, to speed up the image acquisition process, these systems are utilized at a significantly lower than nominal resolution. In SDOCT imaging, practical resolution in the azimuthal axis corresponds to the number of B-scans sampled at relatively equal distances in a *volumetric* scanning scheme (Fig. 1.5). Note that, valid quantitative measurements of retinal disease biomarkers (e.g. drusen [56] volume) are only feasible from B-Scans with known azimuthal displacement.

On the quest to gather useful information from SDOCT through improving the hardware design, one quickly runs into the problem of diminishing returns. Specifically, the optical components necessary to capture very high-quality, dense scans become prohibitively expensive or too sensitive for many practical applications. Unlike alternative approaches that require expensive hardware such as eye tracking systems [22], we propose a software-based image processing solution in this section based on our earlier work [12], that is applicable to virtually any SDOCT imaging system, including the handheld SDOCT systems which are more prone to motion errors [5, 6, 47].

In this section, we introduce a novel application of the super-resolution





FIGURE 1.5

A volumetric SDOCT scan set is a collection of azimuthally sampled B-Scans (a), creating a 3-D view of the retina (b). The braces in (a,b) mark the lower retinal slab shown in (c) containing the shadows from the overlying larger vessels. (d) is the summed voxel projection (SVP) created by axial projection of the lower half of the B-Scans, demonstrating the vessel pattern [22, 28]. Each B-Scan corresponds to one line on the SVP.

framework for improving the azimuthal resolution of SDOCT images. We propose a method based on capturing several repeated fast sparse 3-D scans, followed by detecting and removing the ones affected by motion artifacts, and finally fusing the artifact-free scans. Our approach to reduce motion artifacts in the 3-D tomographic structure, in spirit, is close to the multi-camera time-space super-resolution [48], MRI inter-slice reconstruction [21], and video synchronization [57] problems. However, the proposed reconstruction algorithms and applications are fundamentally different and novel.

1.3.2.1 Proposed Method: Sparse Repeated Imaging

Our goal is to transcend the limitations of SDOCT imaging systems, reduce motion artifacts (Fig. 1.6), and obtain densely sampled, high-quality, and

**FIGURE 1.6**

SVPs from densely sampled volumetric SDOCT scans of two subjects. The red ellipsoids mark motion artifact locations.

accurate 3-D SDOCT images of unpredictably moving structures such as a human eye. In typical SDOCT ophthalmic imaging practice, the region of interest is swept with a relatively high number of B-Scans. For many patients, due to the multiple seconds required to capture scans, a dense scanning strategy is prone to motion artifacts such as blinking.

Alternatively, we propose to capture several (N) sparsely sampled volumetric scans with a significantly lower number of B-Scans than the target resolution. Since the number of frames in each sequence (K) is relatively small, each scan is captured very fast. Therefore, it is reasonable to assume that some of these sequences will be less affected by the abrupt patient motion. We detect such sequences, reorder and interlace their frames, and create a densely sampled, artifact-free representation of the retina. Figure 1.7 represents the main idea, where two sparsely sampled sequences are fused together creating a dense representation of the underlying pathological structures.

Putting together the frames of different scan sets (interlacing) in a correct order is a challenging task. A naive approach involves sorting via pair-wise registration and computing a closeness measure (e.g. normalized cross-correlation or sum-of-squared difference) of all frames. In the case of fusing only two volumetric B-Scan sets, each frame in the first volume sequence is registered to all frames of the second sequence. Then, in the fused output sequence, this individual frame is inserted into the second sequence nearest to the frame in the second sequence with the highest cross-correlation value. This process would be repeated for all remaining frames in the first sequence. Of course, this is a simplified variation of the video synchronization problem, discussed in detail in the computer vision literature [57]. However, aside from the prohibitively heavy computational load of registering large SDOCT data sets, the

SNR of the SDOCT images is significantly lower than the commercial camcorders for which the method in [57] is developed. Therefore, the commonly used closeness measures such as normalized cross-correlation may not always be sensitive enough to discriminate between very small structural changes in the neighboring SDOCT ophthalmic scans (Fig. 1.7).




FIGURE 1.7

Fusing (interlacing) multiple sparsely sampled scan sequences to create an azimuthally higher resolution volume of B-Scans. Indeed, unlike this schematic example, in clinical applications the displacement between sequences might be non-integer as it is induced by patient motion.

To reduce the computational complexity of 3-D registration and improve accuracy, we introduce an alternative global solution based on 2-D registration. Note that, the azimuthal axis displacement is the only motion that we need to estimate to be able to interlace the 3-D volumetric scans. A quick consultation with Figure 1.5 shows that the y -axis in the 2-D SVP images corresponds to the azimuthal axis in the 3-D data volumes. Therefore, instead of dealing with full 3-D datasets, we axially project the input 3-D sequences and create corresponding SVP images (Fig. 1.5). This will reduce the task of registering K sets each with B images (B-Scans) of size $[B \times L]$ pixels, to registering only K images (SVPs) each of size $[B \times B]$ pixels. In essence, we are projecting down into the SCP domain to create a collection of K images that are undersampled in only one dimension (e.g. $D_1 = 1$).

Aside from a significant reduction in data volume, the axial projection reduces the noise in the SVP images by averaging over hundreds of pixels. As SNR of the SVP images is relatively higher, outlier (motion artifact corrupted) image sets can be more accurately detected and excluded from the data pool.

We recover the order of the frames in the dense 3-D output by registering the remaining SVPs. As explained in the next subsection, we calculate the y -axis motion between different SVPs and associate this to the azimuthal motion

parameters (frame number) of the 3-D volumes. For example, an estimated five pixel displacement for the SVPs of two scan sets in y -axis, indicates an offset of five frames in the corresponding B-Scan sequences. Moreover, by estimating the x -axis motion of the SVPs, we recover the lateral registration parameters needed for aligning the fused (interlaced) B-Scans in the final fused 3-D volume.

1.3.2.2 Multiframe Joint Registration

In many super-resolution applications, fast pairwise image registration is sufficient for estimating the relative shifts between the sets of low-resolution images. The basis for this approach is based on the following approximation of Equation (1.5)

$$\mathbf{y}_k = \mathbf{D}\mathbf{S}(\mathbf{v}_k)\mathbf{z} \approx \mathbf{S}(\mathbf{v}'_k)\mathbf{p} + \mathbf{e}_k + \mathbf{a}_k \quad (1.16)$$

where $\mathbf{v}'_k = [v_{1,k}v_{2,k}/D_2]$ is the apparent motion in the undersampled image, \mathbf{p} is the approximate non-aliased portion of the low-resolution image, and \mathbf{a}_k is the aliasing artifacts which we approximately treated as noise. From the simplified model of Equation (1.16), we see that the relationship between a pair of low-resolution images \mathbf{y}_k and \mathbf{y}_j is approximately given by

$$\mathbf{y}_k = \mathbf{S}(\mathbf{v}'_j - \mathbf{v}'_k)\mathbf{y}_j \approx \mathbf{y}_j + (\mathbf{v}'_j - \mathbf{v}'_k)\nabla\mathbf{S}(0)\mathbf{y}_j \quad (1.17)$$

where the second half of the equation is based on the first order Taylor approximation of the shift operator $\mathbf{S}(\mathbf{v})$. In practice, the operators $\nabla\mathbf{S}(0) = [\mathbf{S}_x(0)\mathbf{S}_y(0)]$ are approximately the x and y gradient operators. Equation (1.17) is known as the optical flow constraint and can be used to estimate the shift between any pair of low-resolution frames. Such an approach works as long as the energy in the aliasing artifacts \mathbf{a}_k are minimal. We use the notation $\mathbf{v}'_{j-k} = \mathbf{v}'_j - \mathbf{v}'_k$.

Due to the sub-Nyquist sampling in the azimuthal direction, the SVPs of the sparse, fast-acquired sequences are aliased in the y -axis, complicating the subpixel motion estimation task. Moreover, small estimation bias in the pairwise SVP registration is magnified to a significant misalignment error when several sequences are fused together. Therefore, to minimize the overall motion estimation error, we exploit global consistency conditions in a multiframe motion estimation framework [14, 19]. This bundle-adjusted, optical flow-based technique relies on the fact that the operator describing the motion between any pair of frames must be the composition of the operators between two other pairs of frames. In effect, by incorporating this prior information in the joint motion estimation framework of [14], we minimize the motion estimation bias while having extremely fast registration by estimating motion entirely in the low-resolution domain.

We overcome the errors associated using a global constraint enforcing the transitivity of the pairwise motion estimates. For example, if we consider three

frames, then the transitivity of the motion estimates requires that

$$\mathbf{v}'_{k-j} = \mathbf{v}'_{l-j} + \mathbf{v}'_{k-l} \quad (1.18)$$

Figure 1.8 schematically describes this motion constraint. In the case of multiple translational motion vectors [14], the above conditions can be described by simple linear equations relating the motion vectors between the frames as




FIGURE 1.8

Global motion consistency conditions that exist for any set of images: the operator describing the motion between any pair of images is the composition of the operators between two other pairs of images: $\mathbf{v}_{i,k} = \mathbf{v}_{i,j} + \mathbf{v}_{j,k}$.

$$\mathbf{U}\mathbf{V} = 0, \quad (1.19)$$

where \mathbf{U} is a $[2(K-1)^2 \times 2K(K-1)]$ consistency condition matrix and \mathbf{V} is a vector collecting the set of unknown motion vectors $\{\mathbf{v}'_k\}$. Each row in the sparse matrix \mathbf{U} has only two or three non-zero (± 1) elements. Motion vectors are estimated by minimizing a global cost function such as

$$\Omega_{\text{OF}}(V) = \sum_{i \neq j}^K \|\mathbf{y}_i - \mathbf{y}_j - \mathbf{v}'_{i-j} \nabla \mathbf{S}(0) \mathbf{y}_j\|_1, \text{ s.t. } \mathbf{U}V = 0, \quad (1.20)$$

We used nonlinear programming (“*fmincon*” function in MATLAB) to minimize this cost function. The above conditions for the more general case of affine motion are defined in [14]. This is a simpler but faster implementation of the general framework described in Section 1.2. Since we do not estimate the high-resolution image jointly with the registration parameters our solution is suboptimal. Indeed, mathematically more rigorous solutions for registering aliased images are also possible [42], which increase computational complexity of the proposed algorithm. However, noting the extremely large SDOCT image sets (hundreds of images of size $[512 \times 1024]$ or larger), the goal of our proposed solution is to be practical for clinical implementation rather than mathematically optimal.

1.3.2.3 Experimental Results

The above registration technique recovers the order and the relative azimuthal distance of B-Scans from different scan sets, which can be exploited to reconstruct a dense 3-D view of the imaged pathological structure. Since misaligned or broken vessels are easily detectable in retinal imaging applications, the vessel pattern as seen on the SVP serves as an efficient qualitative measure of the success and accuracy of the overall algorithm. Therefore, we use the estimated motion parameters to reconstruct a fused (super-resolved) 2-D SVP map of retinal vessel structure.

Figure 1.6 shows a dense scanning of a subject, whose motion artifacts have resulted in an SVP with broken vessel structure. From the same subject, we captured 12 sparsely sampled volumetric scans (each with 50 B-Scans) and adjusted the baseline of each image using the fast registration StackReg plug-in (Biomedical Imaging Group; Swiss Federal Institute of Technology Lausanne) [54] for ImageJ (freeware; National Institutes of Health; Bethesda, MD). Following [28], by summing the lower half of the B-Scans in the axial direction, we created SVPs with distinct vessel patterns. After contrast adjustment, four of the six sequences with the highest SVP normalized cross-correlation values were manually selected to be registered. Figures 1.9(a) and 1.9(b) show two corresponding SVPs of these four sequences. Registered and sequentially ordered AVI movies of these four input sequences are available in http://www.duke.edu/~sf59/datasets/SDOCT_SR.avi, screenshots of which are shown in Figure 1.9(f).

We used the multiframe projective bundle-adjusted motion estimation method of [14] to recover the subpixel translational motion parameters of these four SVPs (Sec. 1.3.2.2). We used the fast zero-order classic kernel regression-based super-resolution algorithm described in [52] to reconstruct the fused SVP. Since there are no aliasing artifacts in the x -axis (lateral direction), the

SVP resolution is only enhanced in the y -axis. The SVP of the fused sequence is shown in Figure 1.9(c), which has more details than any of the input SVPs.

As a point of reference, we also captured a gold-standard sequence shown in Figure 1.9(d), which is the visually best of 4 densely sampled volumetric scans (100 B-Scans). The reconstruction accuracy and quality improvement is confirmed by comparing the input and output SVPs to the gold-standard. We believe the small jaggedness in reconstructed vessels of Figure 1.9(c) is mainly due to the dynamic structural deformation of the vessels during the cardiac cycle. Overall, the vessel pattern in Figure 1.9(c) shows fewer discontinuities (blue ellipsoids) compared to the input SVPs. Moreover, due to the less aggressive interpolation in the azimuthal (y -axis) direction, the vessel thicknesses are more accurate in Figure 1.9(c) than in any of the input frames (red ellipsoids).

1.4 Conclusion

We have provided a proof of concept for the applicability of image processing based algorithms as an alternative to expensive hardware for creating robust high quality X-ray mammography and SDOCT ophthalmic imaging. The proposed super-resolution-based algorithm enables ophthalmic imaging practitioners and radiologists to optimally utilize the SDOCT and X-ray systems in their highest resolution capacity.

For the SDOCT case, several implementation variations for improving the efficiency are possible. For example, rather than discarding a whole defected sequence, we may discard only those B-Scans affected by abrupt motion artifacts (e.g. blinking), and use the remaining uncorrupted B-Scans. To produce more visually appealing SVPs, more efficient super-resolution techniques such as the steering adaptive kernel [52] or robust super-resolution [16] may be also exploited. Moreover, incorporation of an advanced adaptive sparse sampling strategy (3-D extension of the method in [13]) in this imaging framework is part of our ongoing research.

While the proposed algorithm efficiently removes abrupt motion artifacts, a practical drawback is the case of imaging objects with constant deformable motion. For example, in the case of imaging pulsing blood vessels, each sparse sequence is associated with a unique position of the blood vessels compared to the background tissues. A possible remedy is synchronizing the start of image acquisition in each sparse sequence with the electrocardiogram (EKG) signal.

We note that two alternative related sparse imaging scenarios can be also considered. One is based on capturing large field of view repeated scans, dense in the azimuthal direction but sparsely sampled in the lateral direction. Then, a classic super-resolution algorithm (e.g. [16, 52]) may reconstruct the lateral resolution of individual B-Scans. Our pilot experimental results have shown

moderate improvements, when imaging objects under a SDOCT microscope. However, in practical clinical trials the difficulty of capturing repeated scans from a unique azimuthal location largely voids the applicability of this strategy.

Alternatively, the authors of the published work in [66] propose to capture azimuthally dense scan sets with a small field of view in the axial-lateral plane (e.g. the *en face* view is divided into four sub-sections, which are imaged subsequently). A customized semi-automatic software stitches the 3-D scan volumes, creating a visually appealing, large field of view, 3-D rendition of the retina. However, unfortunately for the same practical imaging problem noted for the aforementioned strategy, as evident in the experimental results in [66], it is extremely hard (if not impossible) to recover unique, large field of view B-Scans without evident registration artifacts.


As for the mammography, we believe the design of future X-ray imaging systems would benefit from a systematic analysis of the resolution and SNR required for mammographic screening and diagnosis. In the future, we will explore the fundamental tradeoffs between radiation exposure, number of frames, and reconstruction performance. Furthermore, we will investigate more sophisticated redundant wavelet techniques such as curvelets [50] or ridgelets [7] which might show even better performance than the proposed multiframe ForWaRD technique. In fact, recent research has shown that use of more sophisticated wavelets can improve the image quality in other medical imaging applications [31].

We believe this novel application of super-resolution can be used as a stepping stone toward many other image fusion based medical imaging system designs, aimed especially at patients with uncontrollable motion, pediatrics, or hand-held probe imaging. While this chapter was focused on X-ray and OCT image enhancement, similar strategies can be exploited for enhancing the quality of some other volumetric medical imaging devices such as ultrasound.

1.5 Acknowledgment

Our MATLAB software implementation of the noted X-ray algorithms is in part based on the ForWaRD software, developed by Dr. Ramesh D. Neelamani of the Digital Signal Processing group at Rice University (available at <http://www.dsp.rice.edu/software/ward.shtml>). We would like to thank the Editor Prof. Peyman Milanfar for collaborating with us in the original multiframe X-ray motion estimation publications. We thank Bradley A. Bower and Yuankai K. Tao for providing invaluable SDOCT data and assistance. This work was supported in part by North Carolina Biotechnology Center Collaborative Funding Grant #2007-CFG-8005 with Bioptigen, the Duke Translational Medicine Institute Subcontract #12 of NIH Grant #5ULT-RR024128-

03, Knights Templar Eye Foundation Inc. Pediatric Ophthalmology Research Grant, Hartwell Foundation Individual Biomedical Research Award.

**FIGURE 1.9**

(a) and (b) are two representative SVPs of four input retinal SDOCT sequences (50 regularly sampled B-Scans each). (c) is the SVP of the fused sequence (200 irregularly sampled B-Scans). (d) is the Gold-Standard SVP which is the best dense (100 regularly sampled B-Scans) out of 4 such sequences (Figure 1.6(b) is an example of dense sampling of the same subject with motion artifacts). (e) and (f) are the screen shots of the AVI movies of registered four input B-Scan sets and the reordered and interlaced output B-Scan set, respectively (2.1 MB).

Bibliography

- [1] Y. Altunbasak, A. Patti, and R. Mersereau. Super-resolution still and video reconstruction from MPEG-coded video. *IEEE Trans. Circuits And Syst. Video Technol.*, 12(4):217–226, Apr. 2002.
- [2] L.D. Alvarez, J. Mateos, R. Molina, and A.K. Katsaggelos. High resolution images from compressed low resolution video: Motion estimation and observable pixels. *International Journal of Imaging Systems and Technology*, 14(2):58–66, October 2004.
- [3] S. Borman and R.L. Stevenson. Super-resolution from image sequences - a review. In *Proc. of the 1998 Midwest Symposium on Circuits and Systems*, volume 5, Apr. 1998.
- [4] S.G. Chang, B. Yu, and M. Vetterli. Adaptive wavelet thresholding for image denoising and compression. *IEEE Transactions on Image Processing*, 9(9):1532–1546, 2000.
- [5] S.H. Chavala, S. Farsiu, R. Maldonado, D.K. Wallace, S.F. Freedman, and C.A. Toth. Insights into Advanced Retinopathy of Prematurity Using Handheld Spectral Domain Optical Coherence Tomography Imaging. *in press, Ophthalmology*, 2009.
- [6] G.T. Chong, S. Farsiu, S.F. Freedman, N. Sarin, A.F. Koreishi, J.A. Izatt, and C.A. Toth. Abnormal foveal morphology in ocular albinism imaged with spectral-domain optical coherence tomography. *Archives of Ophthalmology*, 127(1):37–44, 2009.
- [7] M. Do and M. Vetterli. The finite ridgelet transform for image representation. *IEEE Trans. Image Process.*, 12(1):16–28, January 2003.
- [8] C. Dunsby and P.M.W. French. Techniques for depth-resolved imaging through turbid media including coherence-gated imaging. *Journal of Physics D: Applied Physics*, 36:207–227, July 2003.
- [9] N.A. El-Yamany and P.E. Papamichalis. Robust color image superresolution: An adaptive m-estimation framework. *EURASIP Journal on Image and Video Processing*, 8:1–12, 2008.
- [10] M. Elad and A. Feuer. Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images. *IEEE Transactions on Image Processing*, 6(12):1646–1658, Dec 1997.

- [11] M. Elad and Y. Hel-Or. A fast super-resolution reconstruction algorithm for pure translational motion and common space invariant blur. *IEEE Transactions on Image Processing*, 10(8):1186–1193, August 2001.
- [12] S. Farsiu, B.A. Bower, J.A. Izatt, and C.A. Toth. Image Fusion Based Resolution Enhancement of Retinal Spectral Domain Optical Coherence Tomography Images. *Invest. Ophthalmol. Vis. Sci.*, 49(5):E–abstract–1845, 2008.
- [13] S. Farsiu, J. Christofferson, B. Eriksson, P. Milanfar, B. Friedlander, A. Shakouri, and R. Nowak. Statistical detection and imaging of objects hidden in turbid media using ballistic photons. *Applied Optics*, 46(23):5805–5822, 2007.
- [14] 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*, pages 1396–1401, July 2005.
- [15] S. Farsiu, D. Robinson, M. Elad, and P. Milanfar. Advances and challenges in super-resolution. *International Journal of Imaging Systems and Technology*, 14:47–57, 2004.
- [16] S. Farsiu, M.D. Robinson, M. Elad, and P. Milanfar. Fast and robust multiframe super resolution. *IEEE Transactions on Image Processing*, 13(10):1327–1344, October 2004.
- [17] A.F. Fercher, W. Drexler, C.K. Hitzenberger, and T. Lasser. Optical coherence tomography - principles and applications. *Reports on Progress in Physics*, 66(2):239–303, 2003.
- [18] G. Golub and V. Pereyra. Separable nonlinear least squares: the variable projection method and its applications. *Institute of Physics, Inverse Problems*, 19:R1–R26, 2002.
- [19] V.M. Govindu. Lie-algebraic averaging for globally consistent motion estimation. In *Proc. of the Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 684–691, July 2004.
- [20] H. Greenspan. Super-resolution in medical imaging. *The Computer Journal*, 52:43–63, 2009.
- [21] H. Greenspan, G. Oz, N. Kiryati, and S. Peled. MRI inter-slice reconstruction using super-resolution. *Magnetic Resonance Imaging*, 20(5):437–446, 2002.
- [22] D. Hammer, R.D. Ferguson, N. Iftimia, T. Ustun, G. Wollstein, H. Ishikawa, M. Gabriele, W. Dilworth, L. Kagemann, and J. Schuman. Advanced scanning methods with tracking optical coherence tomography. *Optics Express*, 13(20):7937–7947, 2005.

- [23] R. Hardie. A fast image super-resolution algorithm using an adaptive wiener filter. *IEEE Transactions on Image Processing*, 16(12):2953–2964, December 2007.
- [24] R. Hardie, K. Barnard, and E. Armstrong. Joint MAP registration and high-resolution image estimation using a sequence of undersampled images. *IEEE Transactions on Image Processing*, 6(12):1621–1633, 1997.
- [25] M.R. Hee, J.A. Izatt, E.A. Swanson, D. Huang, J.S. Schuman, C.P. Lin, C.A. Puliafito, and J.G. Fujimoto. Optical coherence tomography of the human retina. *Archives of Ophthalmology*, 113(3):325–332, 1995.
- [26] D.B. Hess, S.G. Asrani, M.G. Bhide, L.B. Enyedi, S.S. Stinnett, and S.F. Freedman. Macular and retinal nerve fiber layer analysis of normal and glaucomatous eyes in children using optical coherence tomography. *American Journal of Ophthalmology*, 139(3):509–517, 2005.
- [27] D. Huang, E.A. Swanson, C.P. Lin, J.S. Schuman, W.G. Stinson, W. Chang, M.R. Hee, T. Flotte, K. Gregory, C.A. Puliafito, et al. Optical coherence tomography. *Science*, 254(5035):1178–1181, 1991.
- [28] S. Jiao, R. Knighton, X. Huang, G. Gregori, and C. Puliafito. Simultaneous acquisition of sectional and fundus ophthalmic images with spectral-domain optical coherence tomography. *Optics Express*, 13(2):444–452, 2005.
- [29] J.A. Kennedy, O. Israel, A. Frenkel, R. Bar-Shalom, and H. Azhari. Super-resolution in PET imaging. *IEEE transactions on medical imaging*, 25(2):137–147, 2006.
- [30] J.A. Kennedy, O. Israel, A. Frenkel, R. Bar-Shalom, and H. Azhari. Improved image fusion in PET/CT using hybrid image reconstruction and super-resolution. *Int. J. Biomed. Imaging*, 46846, 2007.
- [31] A. Khare and U.S. Tiwary. A new method for deblurring and denoising of medical images using complex wavelet transform. *Proc. IEEE Conference Engineering in Medicine and Biology*, pages 1897–1900, Sept. 2005.
- [32] S. Lertrattanapanich and N.K. Bose. High resolution image formation from low resolution frames using Delaunay triangulation. *IEEE Trans. Image Processing*, 11(12):1427–1441, Dec. 2002.
- [33] N. Nassif, B. Cense, B.H. Park, S.H. Yun, T.C. Chen, B.E. Bouma, G.J. Tearney, and J.F. de Boer. In vivo human retinal imaging by ultrahigh-speed spectral domain optical coherence tomography. *Optics Letters*, 29(5):480–482, 2004.
- [34] R. Neelamani, H. Choi, and R. Baraniuk. Forward: Fourier-wavelet regularized deconvolution for ill-conditioned systems. *IEEE Transactions on Image Processing*, 52(2):418–433, February 2004.

- [35] S.C. Park, M.K. Park, and M.G. Kang. Super-resolution image reconstruction: a technical overview. *Signal Processing Magazine*, 20(3):21–36, 2003.
- [36] V. Patanavijit and S. Jitapunkul. A lorentzian stochastic estimation for a robust iterative multiframe super-resolution reconstruction with Lorentzian-Tikhonov regularization. *EURASIP Journal on Image and Video Processing*, 2007(2), 2007.
- [37] R.R. Peeters, P. Kornprobst, M. Nikolova, S. Sunaert, T. Vieville, G. Mallandain, R. Deriche, O. Faugeras, M. Ng, and P. Van Hecke. The use of super-resolution techniques to reduce slice thickness in functional MRI. *International Journal of Imaging Systems and Technology*, 14(3):131–138, 2004.
- [38] S. Peled and Y. Yeshurun. Superresolution in MRI: Application to human white matter fiber tract visualization by diffusion tensor imaging. *Magnetic resonance in medicine*, 45(1):29–35, 2001.
- [39] L.C. Pickup, D.P. Capel, S.J. Roberts, and A. Zisserman. Bayesian methods for image super-resolution. *The Computer Journal*, 2007.
- [40] C.G. Pieroni, A.J. Witkin, T.H. Ko, J.G. Fujimoto, A. Chan, J.S. Schuman, H. Ishikawa, E. Reichel, and J.S. Duker. Ultrahigh resolution optical coherence tomography in non-exudative age related macular degeneration. *British Medical Journal*, 90(2):191–197, 2006.
- [41] D. Robinson, S. Farsiu, J.Y. Lo, P. Milanfar, and C.A. Toth. Efficient multiframe registration of aliased x-ray images. *Proceedings of the 41th Asilomar Conference on Signals, Systems, and Computers*, pages 215–219, November 2007.
- [42] D. Robinson, S. Farsiu, and P. Milanfar. Optimal registration of aliased images using variable projection with applications to super-resolution. *The Computer Journal*, 52(1):31–42, January 2009.
- [43] D. Robinson and P. Milanfar. Statistical performance analysis of super-resolution. *IEEE Transactions on Image Processing*, 15(6):1413–1428, June 2006.
- [44] M.D. Robinson, C.A. Toth, J.Y. Lo, and S. Farsiu. Efficient fourier-wavelet super-resolution with applications in low-dosage digital x-ray imaging. *submitted to IEEE Transactions on Image Processing*, January 2009.
- [45] R.R. Schultz and R.L. Stevenson. Extraction of high-resolution frames from video sequences. *IEEE Transactions on Image Processing*, 5(6):996–1011, June 1996.

- [46] S.G. Schuman, A.F. Koreishi, S. Farsiu, S. Jung, J.A. Izatt, and C.A. Toth. Photoreceptor layer thinning over drusen in eyes with age-related macular degeneration imaged in vivo with spectral-domain optical coherence tomography. *Ophthalmology*, 116(3):488–496, 2009.
- [47] A.W. Scott, S. Farsiu, L.B. Enyedi, D.K. Wallace, and C.A. Toth. Imaging the infant retina with a hand-held spectral-domain optical coherence tomography device. *American Journal of Ophthalmology*, 147(2):364–373, 2009.
- [48] E. Shechtman, Y. Caspi, and M. Irani. Space-time super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(4):531–545, 2005.
- [49] V.J. Srinivasan, M. Wojtkowski, A.J. Witkin, J.S. Duker, T.H. Ko, M. Carvalho, J.S. Schuman, A. Kowalczyk, and J.G. Fujimoto. High-definition and 3-dimensional imaging of macular pathologies with high-speed ultrahigh-resolution optical coherence tomography. *Ophthalmology*, 113(11):2054–2054, 2006.
- [50] J. Starck, E. Candes, and D. Donoho. The curvelet transform for image denoising. *IEEE Trans. Image Process.*, 11(6):670–684, June 2002.
- [51] M. Stopa, B.A. Bower, E. Davies, J.A. Izatt, and C.A. Toth. Correlation of pathologic features in spectral domain optical coherence tomography with conventional retinal studies. *Retina*, 28(2):298–308, 2008.
- [52] H. Takeda, S. Farsiu, and P. Milanfar. Kernel regression for image processing and reconstruction. *IEEE Trans. Image Process.*, 16(2):349–366, Feb. 2007.
- [53] H. Takeda, S. Farsiu, and P. Milanfar. Deblurring using regularized locally adaptive kernel regression. *IEEE Transactions on Image Processing*, 17(4):550–563, April 2008.
- [54] P. Thevenaz, U.E. Ruttimann, and M. Unser. A pyramid approach to subpixel registration based on intensity. *IEEE Transactions on Image Processing*, 7(1):27–41, 1998.
- [55] C.A. Toth, R. Birngruber, S.A. Boppart, M.R. Hee, J.G. Fujimoto, C.D. DiCarlo, E.A. Swanson, C.P. Cain, D.G. Narayan, G.D. Noojin, et al. Argon laser retinal lesions evaluated in vivo by optical coherence tomography. *American Journal of Ophthalmology*, 123(2):188–98, 1997.
- [56] C.A. Toth, S. Farsiu, A.A. Khanifar, and G.T. Chong. Optical coherence tomography in age-related macular degeneration. In Gabriel Coscas, editor, *Application of Spectral Domain OCT in AMD*, pages 15–34. Springer Medizin Verlag Heidelberg, 2009.

- [57] M. Ushizaki, T. Okatani, and K. Deguchi. Video synchronization based on co-occurrence of appearance changes in video sequences. *Pattern Recognition, International Conference on*, 3:71–74, 2006.
- [58] P. Vandewalle, L. Sbaiz, J. Vandewalle, and M. Vetterli. Super-resolution from unregistered and totally aliased signals using subspace methods. *IEEE Transactions on Signal Processing*, 55(7):3687–3703, July 2007.
- [59] P. Vandewalle, S. Susstrunk, and M. Vetterli. A frequency domain approach to registration of aliased images with application to super-resolution. *EURASIP Journal on Applied Signal Processing*, page Article ID 71459, 2006.
- [60] M. Wojtkowski, R. Leitgeb, A. Kowalczyk, T. Bajraszewski, and A.F. Fercher. In vivo human retinal imaging by fourier domain optical coherence tomography. *Journal of Biomedical Optics*, 7(3):457–463, 2002.
- [61] M. Wojtkowski, V. Srinivasan, T. Ko, J. Fujimoto, A. Kowalczyk, and J. Duker. Ultrahigh-resolution, high-speed, fourier domain optical coherence tomography and methods for dispersion compensation. *Optics Express*, 12(11):2404–2422, 2004.
- [62] N. Woods, N. Galatsanos, and A. Katsaggelos. Stochastic methods for joint registration, restoration, and interpolation of multiple undersampled images. *IEEE Transactions on Image Processing*, 15(1):201–213, January 2006.
- [63] Y. Yasuno, J. Sugisaka, Y. Sando, Y. Nakamura, S. Makita, M. Itoh, and T. Yatagai. Non-iterative numerical method for laterally superresolving Fourier domain optical coherence tomography. *Optics Express*, 14(3):1006–1020, 2006.
- [64] S. Yun, G. Tearney, J. de Boer, N. Iftimia, and B. Bouma. High-speed optical frequency-domain imaging. *Optics Express*, 11(22):2953–2963, 2003.
- [65] R. Zawadzki, S. Jones, S. Olivier, M. Zhao, B. Bower, J. Izatt, S. Choi, S. Laut, and J. Werner. Adaptive-optics optical coherence tomography for high-resolution and high-speed 3d retinal in vivo imaging. *Optics Express*, 13(21):8532–8546, 2005.
- [66] R.J. Zawadzki, A.R. Fuller, S.S. Choi, D.F. Wiley, B.Hamann, and J.S. Werner. Improved representation of retinal data acquired with volumetric fd-oct: co-registration, visualization, and reconstruction of a large field of view. volume 6844, page 68440C. SPIE, 2008.