

High-Resolution Iris Image Reconstruction from Low-Resolution Imagery

R. Barnard^a, V. P. Pauca^a, T. C. Torgersen^a, R. J. Plemmons^{a,c}, S. Prasad^g, J. van der
Gracht^d, J. Nagy^b, J. Chung^b, G. Behrmann^e, S. Mathews^f, M. Mirotznic^f

^a Department of Computer Science, Wake Forest University, Winston-Salem, NC 27109

^b Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322

^c Department of Mathematics, Wake Forest University, Winston-Salem, NC 27109

^d Holospex, Inc., 6470 Freetown Rd., Suite 200-104, Columbia, MD 21044

^e EM Photonics Inc., 51 E. Main Street, Suite 203, Newark, DE 19711

^f Department of Electrical Engineering & Computer Science, Catholic University of America,
Washington, DC 20064

^g Center for Advanced Studies and Department of Physics and Astronomy, University of New
Mexico, Albuquerque, NM 87131

ABSTRACT

We investigate the use of a novel multi-lens imaging system in the context of biometric identification, and more specifically, for iris recognition. Multi-lenslet cameras offer a number of significant advantages over standard single-lens camera systems, including thin form-factor and wide angle of view. By using appropriate lenslet spacing relative to the detector pixel pitch, the resulting ensemble of images implicitly contains subject information at higher spatial frequencies than those present in a single image. Additionally, a multi-lenslet approach enables the use of observational diversity, including phase, polarization, neutral density, and wavelength diversities. For example, post-processing multiple observations taken with differing neutral density filters yields an image having an extended dynamic range. Our research group has developed several multi-lens camera prototypes for the investigation of such diversities.

In this paper, we present techniques for computing a high-resolution reconstructed image from an ensemble of low-resolution images containing sub-pixel level displacements. The quality of a reconstructed image is measured by computing the Hamming distance between the Daugman⁴ iris code of a conventional reference iris image, and the iris code of a corresponding reconstructed image. We present numerical results concerning the effect of noise and defocus blur in the reconstruction process using simulated data and report preliminary work on the reconstruction of actual iris data obtained with our camera prototypes.

1. INTRODUCTION AND MOTIVATION

Recently, a number of research efforts have investigated the potential of multi-lenslet camera systems, notably, the TOMBO project⁷ and MONTAGE projects.¹ The camera system and techniques described in this paper are guided by a series of workshops hosted at the Oak Ridge National Laboratories (ORNL) held during mid-2005 through 2006 as part of an ARDA seedling challenge project. The workshop series entitled “Practical Enhanced-Resolution Integrated Optical-Digital Imaging Camera” (PERIODIC) seeks to design and prototype a multi-lenslet camera which can exploit the observational diversity inherent in a multi-lenslet system. A primary motivation for this project is the need for compact integrated imaging devices capable of resolving the details of a human iris for biometric identification.

Multi-lenslet camera systems enable the use of many types of observational diversities, including phase, polarization, neutral density, and wavelength diversities. For example, post-processing multiple observations made with differing neutral density filters can yield an image having an extended dynamic range. Applications include surveillance, license plate identification, and other situations where natural lighting may cause excessive contrast. Phase diversity, and the use of phase encoding masks offer the potential for extended depth of field.

Under appropriate conditions, an ensemble of images formed by a multiple lenslet camera will implicitly contain subject information at higher spatial frequencies than those present in a single image.¹² One notable exception, is the degenerate case in which all lenslet images are shifted by an integer number of pixels from each other. By choosing an appropriate lenslet spacing, and detector pitch, such degenerate observing conditions can be avoided for a specific camera application, and the anticipated subject to lens distance.

In this paper, we primarily discuss methods to reconstruct a single high-resolution image from an ensemble of (low-resolution) images. Reconstructions are computed using both simulated observations, and using data obtained from the PERIODIC prototype camera. Our goal is to explore systems which can process acquired data into a form suitable for an automated biometric (iris) identification system.

2. PERIODIC IMAGING SYSTEM OVERVIEW

Figure 1 illustrates biometric (iris) data acquisition using a PERIODIC camera system. A lenslet array forms multiple images of a subject iris on a single detector. Under nearly all observing conditions, the multiple images focused on the detector will be sampled with slightly different alignment of the image with respect to the detector pixels. The ensemble of sub-pixel shifted images encodes subject information at higher spatial frequencies than present in each image individually.

The reconstruction of a single high-resolution image from an ensemble of low resolution images proceeds by first estimating the sub-pixel shifts for each pair of low resolution images. A pair of low-resolution images are up-sampled by a feasible resolution factor and an optimal registration of the two images is computed. Here, a “feasible” gain in resolution is closely related to the number of lenslets, as well as lenslet placement.¹² After estimating the sub-pixel shifts, a high resolution image is computed using techniques discussed in Sections §3 and §4. The high resolution iris image may then be used directly in a biometric application.

In practice, image quality from lenslet to lenslet may vary considerably due to fabrication issues. Experience with early prototype cameras suggests that the inclusion of poor quality, low-resolution images into the reconstruction step results in a lower quality final result. Several low resolution images may be excluded from the reconstruction step, and a good result can still be obtained. However, if a system is formed using too few low-resolution images relative to the desired resolution gain factor, the resulting system will be excessively under determined.

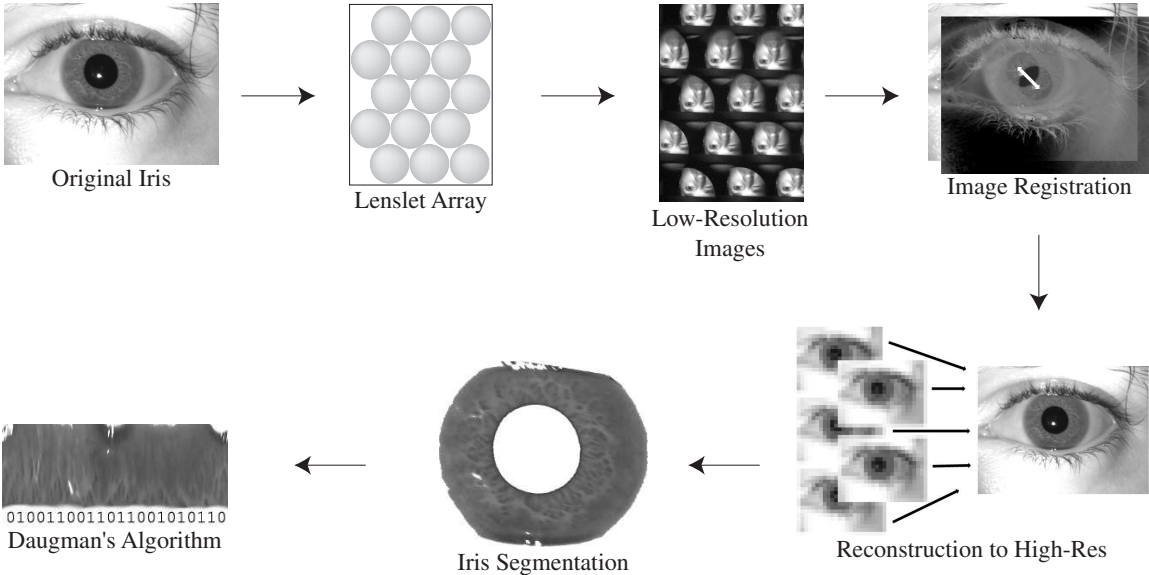


Figure 1. A visual overview of the reconstruction process.

3. IMAGING SYSTEM MODEL

3.1. The Forward Model

The image formation process is naturally described in the terms of continuous functions. However, for brevity, we begin with the discrete approximation to the continuous model. For a multi-lenslet camera, the low-resolution image, g_j formed by the j^{th} lenslet is given by:

$$g_j = DH_j S_j f + \eta_j, \quad (1)$$

where,

- f is a target high-resolution representation for the object, stacked as an $n^2 \times 1$ vector of pixel values.
- S_j is an $n^2 \times n^2$ matrix representing the translation of image f due to the relative position of lenslet j with respect to the object and the detector,
- H_j is an $n^2 \times n^2$ matrix blurring operator associated with lenslet j ,
- D is an $(n/\ell)^2 \times n^2$ decimation matrix operator which represents the transformation from the target $n \times n$ high resolution to the $(n/\ell) \times (n/\ell)$ (low) resolution of the detector,
- g_j is the low-resolution image associated with lenslet j stacked as an $(n/\ell)^2 \times 1$ vector and,
- η_j describes a noise process associated with the j^{th} image.

We refer to ℓ as the *resolution improvement factor*.

Equation (1) serves as a framework for both camera simulation and image restoration. For simulation purposes, we adopt a widely-used noise model⁶ given by:

$$\eta_j = \sqrt{\hat{g}_j} \eta_{j1} + \sigma \eta_{j2} \quad (2)$$

where $\hat{g}_j = DH_j S_j f$. Both η_{j1} and η_{j2} are standard normally distributed random variables with mean equal to zero and variance equal to one. The parameter σ is the standard deviation of the image-independent noise term. The normally distributed term $\sqrt{\hat{g}_j} \eta_{j1}$ in equation (2) approximates the Poisson noise associated with the light detection process. For purposes of image reconstruction, equation (1) leads to a least squares inverse problem whose solution is the desired high resolution reconstruction f .

3.2. The Inverse Problem

Recently, a well-developed theory for digital super-resolution has begun to appear in the literature. See, for example Ref.^{2, 8, 10, 12} The implication for the PERIODIC project is that an understanding of the full potential of multi-lenslet cameras, as well as their inherent limitations begins to emerge. In practical applications, additional challenges arise because of real-world deviations from the models on which the theory is developed. In this section, we consider the inverse problem associated with equation (1) and some of the challenges associated with practical application. For an m -lenslet camera we formulate a least-squares cost functional:

$$J(f) = \left\| \begin{bmatrix} g_1 \\ g_2 \\ \vdots \\ g_m \end{bmatrix} - \begin{bmatrix} DH_1 S_1 \\ DH_2 S_2 \\ \vdots \\ DH_m S_m \end{bmatrix} f \right\|_2^2 \quad (3)$$

whose minimum gives the desired estimate \hat{f} for the high resolution image f , i.e.,

$$\hat{f} = \operatorname{argmin}_f \{ J(f) \} \quad (4)$$

If the number of lenslets m and the resolution improvement factor ℓ satisfy $m = \ell^2$, then the block matrix in equation (3) is square, and the system is fully determined. In practice, we find often that $m < \ell^2$ and the system is under-determined. For values of m reasonably close to ℓ^2 , acceptable reconstructions can still be obtained. Smaller values of m also reduce the computational cost of the reconstruction.

Equations (3) and (4) also present several variations or levels of difficulty, depending on the extent of a priori knowledge. Several relevant issues are discussed in subsections §3.3 to §3.5 below.

3.3. Estimating the PSF, H_j

In practice, the point spread function H_j is not easily estimated, since the detector only records low-resolution images, while H_j operates on the high resolution unknown, f . If S_j is constrained to shifts corresponding to an integral number of high resolution pixels, then, H_j and S_j commute in equation (1). However, the decimation matrix D does not commute with H_j , thus, a restoration step using a low-resolution representation of the point spread function, does not yield information which could be directly helpful for solving the inverse problem in equation (4) for f .

In special cases, H_j may be known. For example, for accurately focused subjects, and sufficiently accurately constructed lenslets, the H_j could be approximated by the diffraction limited PSF, given knowledge of the optical system parameters, such as focal length, F-number, and wavelength. Unfortunately, even simple focus errors are difficult to know accurately, since focus is affected by small manufacturing errors in the camera and by subject position.

Further, low-cost, commodity, lenslets often suffer from large aberrations, including field curvature and other effects which lead to spatially-varying blurs. While much work⁹ has been done on the subject of digitally restoring images containing spatially varying blur, the computational cost of such methods are considered prohibitive for the practical applications of the PERIODIC camera. Phase encoding masks¹³ and corresponding (low cost) post-processing are considered to be promising future directions to overcome spatially varying blur, as well as provide extended depth of field. The PSF of a phase encoding mask can be expected to be known with sufficient accuracy to ensure a good reconstruction.

3.4. Estimating the Sub-pixel Shift Matrix, S_j

For objects at infinity, sub-pixel shifts are due to lenslet placement relative to the detector pixel pitch. In practice, the sub-pixel shifts may be measured experimentally. Using appropriate equipment and careful technique, reasonably accurate estimates can be expected. For nearby objects, parallax implies that the sub-pixel shifts are dependent on the subject-to-lens distance. Thus, for general purpose applications, the sub-pixel shifts must be estimated from the collection of low-resolution images for each new high-resolution image produced.

Here, two possible scenarios may be pursued. One can consider the *joint estimation problem* in which both S_j and f in equation (1) are sought simultaneously in a single joint-optimization problem. Several promising approaches have been investigated.³

For the PERIODIC camera described in this paper, a sequential two-step approach is chosen. Our approach is motivated by a specific imaging task, i.e., biometric (iris) identification, including the need to process imperfect “real-world” data with a response time acceptable to the end user. Because of the relative rigidity of the detector and the lenslet array, sub-pixel shifts can be assumed to be uniform across the image, and limited to translations, and possibly, small-angle rotations. In the prototype camera, adjustable stages ensure that the angle of rotation is sufficiently close to zero so that it may be ignored. Thus, only rigid translations remain. In the first phase of computation, the sub-pixel shift operator S_j is estimated from the low-resolution images by cross-correlation and enumeration. During the second phase, each S_j is held fixed, and the least squares problem described in equations (3) and (4) is solved using the familiar conjugate gradient method (CGLS) to obtain the desired high-resolution image f .

3.5. Estimating the Decimation Matrix, D

The decimation matrix is perhaps the easiest to estimate of the three matrix factors, since the detector geometry is usually known or is at least measurable. In the simplest case, the sub-sampling process can be approximated by integrating the high resolution pixels over an appropriate region. More sophisticated approaches might consider the effects of inter-pixel gap, and the shape and layout pattern of the active area of the detector pixels. Accurately representing the active area for each low-resolution pixel is especially a concern for some types of CMOS detectors.

4. IMAGE REGISTRATION AND HIGH RESOLUTION RECONSTRUCTION

As previously mentioned, a sequential two-step approach is chosen to minimize equation (3) with respect to both the sub-pixel shift operators S_j and the high resolution unknown f . More specifically, given the low resolution images g_j a registration problem is solved to yield S_j . Then the minimization problem in equation (4) is solved for f iteratively using the well-known conjugate gradients method.

4.1. Sub-Pixel Image Registration

Image registration problems arise in a wide variety of applications including astronomy, remote sensing, and medical imaging. The problem may be described as the process of transforming different sets of data into a single coordinate system, to enable the comparison or integration of data obtained from different measurements. While extensive literature on the topic has been published over the last several decades, see e.g. Ref.,¹⁵ performance of registration methods remains largely application data dependent. We have explored two registration methods for the purpose of rapidly estimating sub-pixel shifts between pairs of PERIODIC images: one based on the normalized cross-correlation method and the second based on the minimization of an appropriate distance or similarity metric.

4.1.1. Approach based on the Normalized Cross-Correlation

The normalized cross-correlation $C(x, y)$ of two real functions $h_1(x, y)$ and $h_2(x, y)$ is defined as:⁷

$$C(x, y) = \frac{\iint h_1(x', y') h_2(x' - x, y' - y) dx' dy'}{\left[\iint |h_1(x, y)|^2 dx dy \right]^{1/2} \left[\iint |h_2(x, y)|^2 dx dy \right]^{1/2}} \quad (5)$$

Assuming h_1 and h_2 have a similar distribution, the relative shift between h_1 and h_2 is given by $(x_*, y_*) = \arg \max_{x, y} C(x, y)$. For sampled images g_0 and g_j , a discretized correlation sum is used to approximate equation (5), yielding values $C(x_i, y_i)$ at discrete points (x_i, y_i) . Interpolation may be used to better estimate the location of the peak.¹⁵

An alternative approach consists of calculating the discrete cross-correlation using up-sampled images. The accuracy of this approach depends on the up-sampling method as well as on the amount of high frequency information content available in the data. For the iris images, our approach is summarized as follows:

- Up-sample the low-resolution images g_0 and g_j by a factor $r \geq \ell$ using either bilinear or bicubic interpolation,
- Filter the up-sampled images using high-pass filtering and thresholding to isolate the specular reflection in the pupil, and
- Calculate the discrete cross-correlation between the up-sampled resulting filtered images and find the displacements corresponding to the maximum correlation.

4.1.2. Approach based on Minimization of a Distance Metric

Minimizing the distance between two different images of the same object in some metric space is a natural approach for image registration. Several distance metrics may be applied, including Euclidean distance, mutual information, and normal gradient fields.

Using a Euclidean distance metric, our image registration problem can be formulated as follows. Let g_0 denote a reference image, and let g_j denote a target image of interest, both represented by stacked $(n/\ell)^2 \times 1$ vectors. Let $T(\Delta x, \Delta y)$ denote a $(n/\ell)^2 \times (n/\ell)^2$ image shift matrix operator whose entries are a continuous function of the shift offsets Δx and Δy . Naturally, non-integer shifts involve some form of interpolation, and a number of widely-used methods (e.g., nearest neighbor, bilinear, etc.) are available to form $T(\Delta x, \Delta y)$. A detailed description of methods for constructing the shift operator from sub-pixel displacements can be found in Ref.³

Then, the registration problem may be described by:

$$(\widehat{\Delta x}, \widehat{\Delta y}) = \operatorname{argmin}_{\Delta x, \Delta y} \left\{ \|g_0 - T(\Delta x, \Delta y)g_j\|_2^2 \right\} \quad (6)$$

Unfortunately, equation (6) is extremely difficult to solve in general, since a textured subject can lead to numerous local minima. A more tractable formulation of the registration problem given by equations (7) and (8) below. The basic idea is to seek an optimal registration over a discrete set of integer displacements using two up-sampled images.

Let $S_j(\Delta p, \Delta q)$ denote an $(rn/\ell)^2 \times (rn/\ell)^2$ matrix representing a shift by integer numbers Δp and Δq of high-resolution pixels. Each shift of one high-resolution pixel corresponds to a sub-pixel shift by $1/r$ low-resolution pixels. Let U denote an $(rn/\ell)^2 \times (n/\ell)^2$ up-sample matrix operator. We define an objective functional K ,

$$K(\Delta p, \Delta q) = \|Ug_0 - S_j(\Delta p, \Delta q)Ug_j\|_2^2 \quad . \quad (7)$$

Then, the registration problem can be formulated as:

$$(\widehat{\Delta p}, \widehat{\Delta q}) = \operatorname{argmin}_{\Delta p, \Delta q} \left\{ K(\Delta p, \Delta q) \right\} \quad (8)$$

Equation (8) is solved by starting with the correlation approach described in Section §4.1.1. The shifts indicated by the peak correlation are then refined by enumeration of all shifts contained in a small neighborhood of the peak and evaluation of the distance metric (7). This approach is computationally viable because in practice, only a small neighborhood is required.

For high accuracy, it is natural to choose a relatively large value of the up-sample factor r . However, as the up-sample factor increases, the corresponding processing time implied by equation (8) grows geometrically. A more efficient algorithm can be constructed by starting with a small up-sample factor, e.g., $r = 2$ which yields an initial estimate of the sub-pixel shifts. The initial estimate can then be recursively improved by up-sampling an appropriate neighborhood containing the current approximate solution, and re-computing the minimal distance metric in the up-sampled region.

4.2. Reconstruction

Once the sub-pixel displacements are computed for all images $g_j, j = 1, 2, \dots, m$, the shift operators S_j in equation (3) are well defined. Reconstruction proceeds by minimizing $J(f)$ by the well known method of conjugate gradients (CGLS). While a number of accelerated-convergence methods, such as pre-conditioned conjugate gradients are well known, we adopt CGLS for comparative purposes as illustrated in Section §5. Section §5 describes a series of numerical experiments, and uses a biometric measure, i.e., the Hamming distance between Daugman⁴ iris codes to evaluate image quality.

5. NUMERICAL RESULTS

In this section, we explore the efficacy of the sequential two-step approach described in Section §4 for the reconstruction of iris imagery. First, we apply our approach to simulated iris data to characterize the effect of noise and misfocus blur on the achievable reconstruction quality. We use both the relative error of the reconstructed image and the Hamming distance⁴ between the reconstructed iris code and a reference iris code as image quality metrics. Second, our approach is applied to real iris imagery obtained from a multi-lenslet PERIODIC prototype camera. Hamming distance (HD) is used to quantify the quality of the reconstructed iris data. Finally, we define a Hamming distance improvement factor $b(\sigma)$ to facilitate a comparison between high-resolution reconstructions and low-resolution images, at the same noise level σ ,

$$b(\sigma) = \frac{\text{HD of reference low-res image}}{\text{HD of reconstructed high-res image}} . \quad (9)$$

5.1. Characterization of Noise and Misfocus Effects on Image Reconstruction

We have conducted a large number of experiments to characterize the effect of various factors in the image reconstruction process, including noise, misfocus, registration and deblurring. We present results concerning the effect of noise and blur in the reconstruction process and comment on the effect of mis-registration and inaccurate deblurring.

The simulated iris data was constructed using the image formation model described in Section 3. The target high-resolution image used to form f is of size 296×296 . Ten low-resolution images of size 74×74 were constructed using equation (1) to form vectors g_j , ($j = 1, 2, \dots, 10$). Thus, the resolution improvement factor is $\ell = 4$. These 10 low-resolution images contain different 1/4-pixel shifts, as characterized by the shifting operators S_j . The blurring operators H_j are given by either the identity matrix I (for the case of no blurring) or by Gaussian blurs with standard deviation equal to 1, 5, and 10. The noise is modeled using equation (2) including both signal-dependent and signal-independent terms⁶ for values of $\sigma = c \cdot \max(g_j)$ for each $c \in \{0.001, 0.01, 0.05\}$.

For the reconstruction of simulated data, we consider the ideal case of known shifting operators S_j and compute the solution of the inverse problem given by equation (4). Notice that the reconstruction matrix $[DH_j S_j]$ is underdetermined ($5/8 * 296^2 \times 296^2$), and thus generally consistent with practical scenarios.

Reconstruction of Simulated Data Containing no Defocus Blur ($H_j = I$). We first consider the case where the object is placed at the optimal object-to-camera distance and hence contains no defocus blur. Figure 2a-c shows the effect of increasing noise on the relative error and Hamming distance at each iteration. The curves labeled as "baseline" indicate performance in the absence of noise. The dotted line at 0.32 is the matching cut-off for HD.* The dash-dot line is the HD of a reference low-resolution image. Later plots are illustrated in a similar manner.

The relative error curves show the semi-convergent behavior typical of iterative algorithms, such as CGLS, due to the reconstruction of the noise subspace. Notice that while the baseline relative error continues to (slowly) decay, the relative error in the presence of noise quickly reaches a minimum and starts to increase. The HD of the reconstructions is however equal to 0 at the first few iterations for all three cases and eventually goes up as noise increases. Thus, in terms of achievable HD, the high-resolution images produce HD=0 for all noise values, whereas the HD of the reference low-resolution image is 0.17 for noise=0.05.

Reconstruction of Simulated Blurred Data without a Priori Knowledge of the Blurring Process.

We now consider the case where the object is misplaced with respect to the camera, producing an image with a proportional amount of blur. We assume no knowledge of the blur in the reconstruction process, i.e. $H_j = I$. Figure 3a-c shows the effect of increasing noise on the relative error and Hamming distance at each iteration.

Clearly, without knowledge of the blurring process, the relative error quickly increases after the first iteration. However the HD is significantly improved with respect to that of the reference low-resolution image over the first

*All tests performed compared images of the same eye, thus true positive matches are expected.

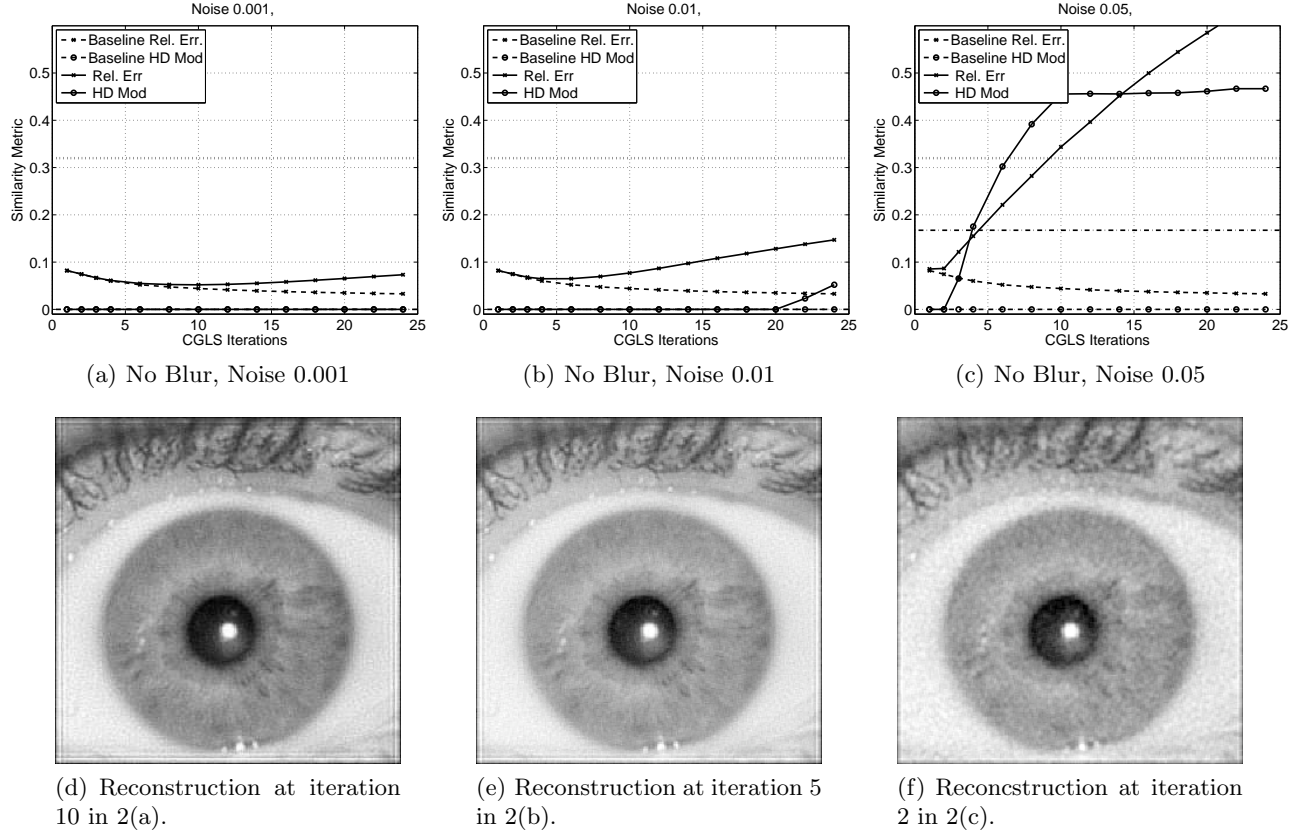


Figure 2. Reconstruction of simulated well-focused low-resolution iris data containing increasing amounts of noise. Dotted line at 0.32: HD matching cutoff. Dash-dot line: HD of reference low-resolution image (noise=0.001: HD(ref)=0, noise=0.01: HD(ref)=0, noise=0.05: HD(ref)=0.17).


few iterations. Thus, in terms of achievable HD, the improvement factors are : $b(\text{noise} = 0.001) = 0.08/0.03 = 2.7$, $b(\text{noise} = 0.01) = 0.17/0.05 = 3.4$, and $b(\text{noise} = 0.05) = 0.33/0.21 = 1.6$.

Reconstruction of Simulated Blurred Data Using a Priori Knowledge of the Blurring Process. We now consider the case where full knowledge of the blurring process can be ascertained prior to and exploited during the reconstruction process. Figure 4a-c shows the effect of increasing noise on the relative error and Hamming distance at each iteration.


As expected, full knowledge of the optical PSF significantly improves both the achievable relative error and HD. Increased amounts of noise can however lessen this improvement as shown in Figure 4c. In terms of achievable HD, the improvement factors are : $b(\text{noise} = 0.001) = 0.8/0.1 = 8$, $b(\text{noise} = 0.01) = 0.17/0.03 = 5.7$, and $b(\text{noise} = 0.05) = 0.33/0.14 = 2.3$. In a realistic scenario these high HD improvement factors would be very difficult to achieve, due to the difficulties associated with estimation of the optical PSF (see Section 3.3).

5.2. Reconstruction of PERIODIC Imagery


In this section we present the results of applying our reconstruction approach to actual iris imagery obtained from a multi-lenslet PERIODIC camera prototype. The camera contains a maximum of 15 lenslets, each about 1.5mm in diameter, with a focal length of about 5mm. It can be used with either a CMOS detector having 3000x2208 pixels with a $3.5\mu\text{m}$ pixel pitch, or with a CCD detector having 1040x1392 pixels with a $6.5\mu\text{m}$ pixel pitch. Representative raw data from this camera is shown in Figure 5.



(d) Reconstruction at iteration 1 in 3(a).



(e) Reconstruction at iteration 1 in 3(b).



(f) Reconstruction at iteration 1 in 3(c).


Figure 3. Reconstruction of simulated defocused low-resolution iris data containing increasing amounts of noise. Dotted line at 0.32: HD matching cutoff. Dash-dot line: HD of reference low-resolution image (noise=0.001: HD(ref)=0.08, noise=0.01: HD(ref)=0.17, noise=0.05: HD(ref)=0.33).

Data obtained with the CMOS detector contains a significant amount of signal-dependent as well as fixed-pattern noise. The latter noise is clearly visible in Figures 7a-b, where it shows as bright specs distributed in a pattern across the image. Raw data obtained with either the CMOS or CCD detectors is appropriately segmented based on the lenslet geometry before reconstruction.


We apply both normalized cross-correlation (NCC) and the distance metric based optimization methods described in Section 4 for registration. Results obtained with NCC were poor; this may have been due in part to dominance of signal-dependent noise in the raw data.⁵ Reconstruction of the iris data using these poor registration values gave essentially no gain in the HD improvement factor. In contrast registration based on the minimization approach of equation (8) produced parameters (lateral offsets Δp and Δq) of much higher accuracy. Figure 6 shows convergence of these parameters with upsample factor r for a pair of CCD low-resolution images.

A summary of the achievable HD values obtained in the reconstruction of both CMOS and CCD iris data is shown in the Table 1. Reconstruction with partial data refers to the use of 5 out of 15 CMOS and 10 CCD low-resolution images in the reconstruction process. The corresponding low-resolution images and high-resolution reconstructions are shown in Figure 7. Notice that for the CMOS detector, the HD for reconstruction with full data (0.21) is higher than the HD for reconstruction with partial data (0.14). This shows that inclusion of low quality frames can negatively impact the reconstruction.


Figure 8 shows the Hamming distance at each iteration of CGLS. The HD value decreased in the first two iterations of CGLS and then increased with subsequent iterations. This is expected, in light of our experimentation




(a) Gaussian blur $\sigma = 5$, Noise 0.001, (b) Gaussian blur $\sigma = 5$, Noise 0.01, (c) Gaussian blur $\sigma = 5$, Noise 0.05, Known Blur



(d) Reconstruction at iteration 6 in 3(a).



(e) Reconstruction at iteration 6 in 3(b).



(f) Reconstruction at iteration 3 in 3(c).

Figure 4. Reconstruction of simulated defocused low-resolution iris data containing increasing amounts of noise. Dotted line at 0.32: HD matching cutoff. Dash-dot line: HD of reference low-resolution image (noise=0.001: HD(ref)=0.08, noise=0.01: HD(ref)=0.17, noise=0.05: HD(ref)=0.33).

	CMOS	Improv. Factor	CCD	Improv. Factor
HD of reference low-res image	0.23	-	0.28	-
HD of high-res reconstruction from partial data	0.14	1.6	0.25	1.12
HD of high-res reconstruction from full data	0.21	1.1	0.19	1.5

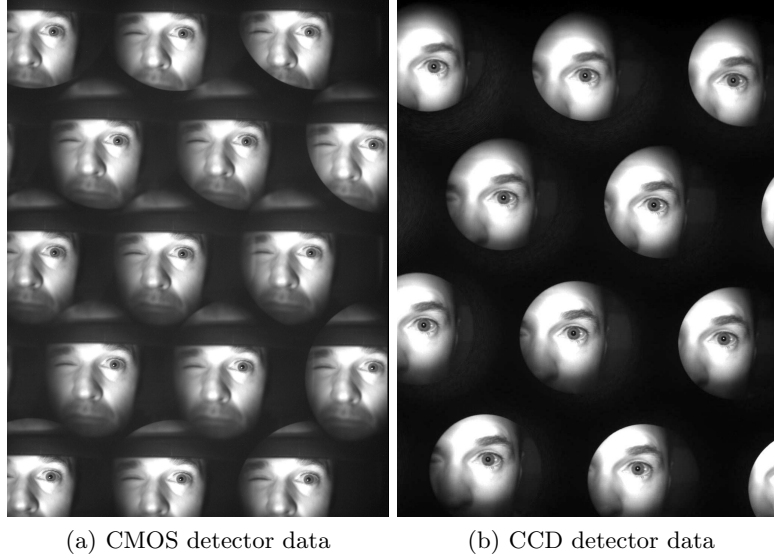
Table 1. Hamming distance value and improvement factor for the CMOS and CCD data.

results with simulated blurred data, since no knowledge of the optical PSFs is employed in the reconstruction process. The maximum HD improvement factors were 1.6 and 1.5 for CMOS and CCD reconstructions, respectively.

6. CONCLUSIONS AND FUTURE WORK

We have demonstrated the viability of a lenslet array based camera for iris recognition using both simulated and actual data. The computational techniques discussed here result in reconstructed images with a consistent improvement over single lenslet images, even with under-determined systems. Simulation results strongly suggest that obtaining a good estimate of the high resolution point spread function for each lenslet is an important challenge, since the quality of reconstruction is sensitive to errors in the PSF estimate.

Naturally, challenges remain, including refinement of methods for estimating the sub-pixel shifts from the data, as well as fabrication challenges to ensure the highest quality and consistency of data across the ensemble



(a) CMOS detector data

(b) CCD detector data

Figure 5. Real data captured using our PERIODIC prototype prototype. The iris is effective 60 pixels across in the CMOS data and 30 pixels across in the CCD data.

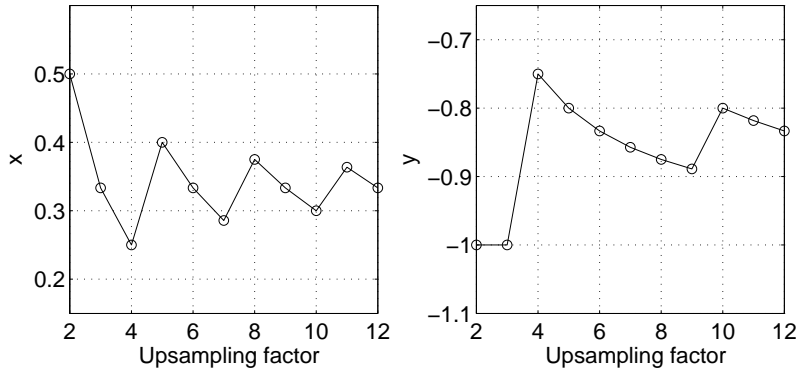


Figure 6. Convergence of Δp and Δq with upsampling factor r .

of images. We wish to further improve the optics and optomechanics so that all sub-images are of good quality. Algorithm performance and system response time are areas of future work, including the possible use of graphics processors (GPUs) for high performance computing with low space and power requirements.

7. ACKNOWLEDGMENTS

The research described in this paper was supported by ARDA under contract number 2364-AR03-A1, for which the authors offer their sincere appreciation. The authors also express their gratitude to the Army Research Office (ARO) for their kind support (award numbers W911NF-05-1-0402 and DAAD19-00-1-0540). We thank Iridian Technologies, Inc. for providing their implementation of the Daugman algorithm under a cooperative agreement with Wake Forest University. Also, we thank Oak Ridge National Laboratory (ORNL) and UT-Battelle for graciously hosting our workshop series in their facility. Finally, we wish to thank Terry Sue Tyrrell of Pfafftown, NC for her assistance in collecting data used in the simulations discussed in Section §5.




Figure 7. Reference images and reconstructions for the CMOS detector (top-row) and CCD detector (bottom-row) data.

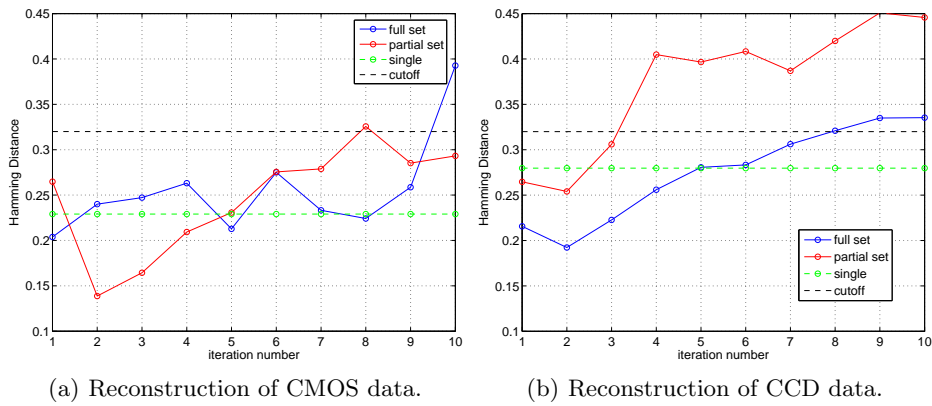


Figure 8. Hamming distance values per iteration of CGLS.

REFERENCES

1. <http://www.disp.duke.edu/projects/imaging>, 2006.
2. BOSE, N. Superresolution from image sequence. In *Proc. 32nd Applied Imagery Pattern Recognition Workshop* (2003), pp. 81–86.
3. CHUNG, J., HABER, E., AND NAGY, J. Numerical methods for coupled super-resolution, 2006. To appear in *Inverse Problems*.

4. DAUGMAN, J. The importance of being random: statistical principles of iris recognition. *Pattern Recognition* 36 (2003), 279–291.
5. GRATADOUR, D., MUGNIER, L., AND ROUAN, D. Sub-pixel image registration with a maximum likelihood estimator. application to the first adaptive optics observations of arp 220 in the l' band. *Astronomy and Astrophysics* 443, 1 (2005), 357–365.
6. JAIN, A. K. *Fundamentals of digital image processing*. Information and System Sciences. Prentice-Hall International, 1989.
7. KITAMURA, Y., SHOGENJI, R., YAMADA, K., MIYATAKE, S., MIYAMOTO, M., MORIMOTO, T., MASAKI, Y., KONDOU, N., MIYAZAKI, D., TANIDA, J., AND ICHIOKA, Y. Reconstruction of a high-resolution image on a compound-eye image-capturing system. *Applied Optics* 43, 8 (2004), 1719–1727.
8. LIN, Z., AND SHUM, H.-Y. Fundamental limits of reconstruction-based superresolution algorithms under local translation. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 26, 1 (2004), 83–97.
9. NAGY, J., AND O'LEARY, D. Restoring images degraded by spatially-variant blur. *SIAM J. Sci. Comput.* 19 (1998), 1063–1082.
10. PARK, S., PARK, M., AND KANG, M. Super-resolution image reconstruction: A technical overview. *IEEE Signal Processing Magazine* 20, 3 (2003), 21–36.
11. PLEMMONS, R., HORVATH, M., LEONHARDT, E., PAUCA, V., PRASAD, S., ROBINSON, S., SETTY, H., TORGERSEN, T., VAN DER GRACHT, J., DOWSKI, E., NARAYANSWAMY, R., AND SILVEIRA, P. Computational imaging systems for iris recognition. In *Proc. SPIE, Advanced Signal Processing Algorithms, Architectures, and Implementations XIV* (Denver, CO, Aug. 2004), vol. 5559, SPIE, SPIE, pp. 346–357.
12. PRASAD, S. Digital superresolution and the generalized sampling theorem, 2006. Submitted to the Optical Society of America.
13. PRASAD, S., PAUCA, V., PLEMMONS, R., TORGERSEN, T., AND VAN DER GRACHT, J. Pupil-phase optimization for extended-focus aberration-corrected imaging systems. In *Proc. SPIE, Advanced Signal Processing Algorithms, Architectures, and Implementations XIV* (Denver, CO, Aug. 2004), vol. 5559, SPIE, SPIE, pp. 335–345.
14. ROCHE, A., MALANDAIN, G., AND AYACHE, N. Unifying maximum likelihood approaches in medical image registration. *Pattern Recognition* 11, 1 (2000), 71–80.
15. ZITOVÁ, B., AND FLUSSER, J. Image registration methods: a survey. *Image and Vision Computing* 21 (2003), 977–1000.