

Penrose Pixels for Super-Resolution

Moshe Ben-Ezra, *Member, IEEE*, Zhouchen Lin, *Senior Member, IEEE*,
Bennett Wilburn, *Member, IEEE*, and Wei Zhang

Abstract—We present a novel approach to reconstruction-based super-resolution that uses aperiodic pixel tilings, such as a Penrose tiling or a biological retina, for improved performance. To this aim, we develop a new variant of the well-known error back projection super-resolution algorithm that makes use of the exact detector model in its back projection operator for better accuracy. Pixels in our model can vary in shape and size, and there may be gaps between adjacent pixels. The algorithm applies equally well to periodic or aperiodic pixel tilings. We present analysis and extensive tests using synthetic and real images to show that our approach using aperiodic layouts substantially outperforms existing reconstruction-based algorithms for regular pixel arrays. We close with a discussion of the feasibility of manufacturing CMOS or CCD chips with pixels arranged in Penrose tilings.

Index Terms—Super-resolution, Penrose tiling, CMOS sensor, CCD sensor.

1 INTRODUCTION

RECENT research in super-resolution (SR) has raised significant concerns regarding the usability of reconstruction-based super-resolution algorithms (RBA [7]) in the real world. Baker and Kanade [7] showed that the condition number of the linear system used in RBA and the volume of solutions grow quickly with increasing magnification. Lin and Shum [25] showed that the effective magnification factor can be at most 5.7. Zhao and Sawhney [35] showed that even properly aligning local patches for SR is difficult. More recent RBA algorithms using statistical models and image priors [34], [30], [20] are more robust and accurate, but the maximum magnification factor remains relatively low (2x to 4x).

Overcoming these limitations requires a new approach to RBA. As noted by Baker and Kanade [7], RBA can be divided into two steps: deblurring optical blur and enhancing resolution. Multiple images taken at small camera displacements provide little or no additional information with respect to the optical blur, so the first step is mostly a blind image deblurring. Moreover, real optical blur is rarely or never shift invariant (and therefore cannot be expressed by a single point spread function) and changes with focus and aperture. This makes the problem of optical deblurring nontrivial at best.

This paper focuses on the second aspect of RBA: detector resolution enhancement using multiple images. Optical deblurring can be applied to the result later, provided that the lens properties are known. There is a significant

technological gap, however, between the theoretical optical resolution limits and current sensor resolutions, particularly for short wavelengths (380-400 μm). This is true for high-quality sensors with large pixels (9 μ to 15 μ) as well as ones with very small pixels (2 μ to 4 μ). Moreover, sensor technology advances more slowly than may be expected [12], while physics is already exploring the feasibility of a “perfect lens” using materials with negative indexes of refraction [29]. Therefore, there is a clear need for resolution enhancement at the sensor level.

1.1 Related Work

Roughly speaking, SR algorithms can be categorized into four classes [11], [27], [15]. Interpolation-based algorithms register low-resolution images (LRIs) with the high-resolution image (HRI), then apply nonuniform interpolation to produce an improved-resolution image, which is then deblurred. Frequency-based algorithms try to de-alias the LRIs using the phase differences between the LRIs. Learning-based algorithms (e.g., [17], [7]) incorporate application-dependent priors to infer the unknown HRI. Reconstruction-based algorithms rely on the relationship between the LRIs and the HRI and assume various kinds of priors on the HRI in order to regularize this ill-posed inverse problem. Among these four categories of algorithms, RBAs are the most commonly used SR algorithms. RBAs usually first form a linear system

$$L = PH + E, \quad (1)$$

where L is the column vector of the irradiances of all the low-resolution pixels (LRPs), H is the vector of the irradiances of the HRI, P gives the weights of the high-resolution pixels (HRPs) in order to obtain the irradiance of the corresponding LRPs, and E is the noise. Past methods to solve (1) for the HRI include maximum a posteriori (MAP) [21], [14], regularized maximum likelihood (ML) [14], projection onto convex sets (POCS) [28], and iterative back projection [23].

In all previous work, the LRPs appear on the left hand side of the system (1) and the LRI pixel layouts are all regular grids, with square pixels. Based on such a

- M. Ben-Ezra and Z. Lin are with Microsoft Research Asia, 5F, Beijing Sigma Center, 49 Zhichun road, Haidian district, Beijing 100190, China. E-mail: {mosheb, zhoulin}@microsoft.com.
- B. Wilburn is with Refocus Imaging, 10607 Vickers Dr., Vienna, VA 22181. E-mail: bennett@stanfordalumni.org.
- W. Zhang is with the Department of Information Engineering, The Chinese University of Hong Kong, Room 703, Ho Sin Hang Engineering Building, The Chinese University of Hong Kong. E-mail: zzw009@ie.cuhk.edu.hk.

Manuscript received 16 July 2009; revised 28 July 2010; accepted 22 Oct. 2010; published online 29 Nov. 2010.

Recommended for acceptance by D. Forsyth.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2009-07-0456.

Digital Object Identifier no. 10.1109/TPAMI.2010.213.

configuration, both the practical and theoretical analyses [7], [25] have shown that the magnification factor is limited to a relatively small number.

Nonrectangular pixel layouts, mostly hexagonal, have been studied [1] for a range of imaging tasks, such as edge detection [26]. Nonrectangular layouts have also been used to increase pixel density, as in the new Fujifilm super CCD. However, these sensors all use periodic tiling and suffer from the same limitations as rectangular tiling with respect to super-resolution. Random sampling was used in the context of compressive sensing [33], [16]. Although Penrose tiling is quasi-random, the compressive sensing framework is very different because RBA is not compressive, and random sampling (tiling) is not considered an aperiodic tiling since the set of tiles is not finite.

1.2 Our Contributions

We increase the RBA magnification factor by breaking the two aforementioned conventions. Rather than using a regular pixel grid for the LRIs, we use *aperiodic* layouts for the detector, resulting in LRIs with an irregular pixel layout. The irregular layout leads to a much more independent equation set. Most importantly, since our layout has no translational symmetry, we can use larger displacements (multiples of half a pixel) between LRIs without having the grid repeat itself. This enables computation of the HRI with larger magnifications. For regular grids, by contrast, the effective displacement is modulo pixel size, which limits the number of different displacements that are at least ϵ apart. We argue that manufacturing a Penrose pixel image sensor is feasible given current technologies and also discuss the potential benefits of such sensors that are not directly related to super-resolution.

To recover a high-resolution image from the raw data captured by the aperiodic pixel layouts, we propose a variant of the traditional error back projection algorithm. Rather than using the LRPs directly for the left-hand side of (1), we upsample the LRIs to the high-resolution grid to match the detector's actual layout (for example, nonsquare pixels with small gaps between them), as shown in Figs. 1a and 1d. For a perfectly square pixel layout, this is identical to a nearest neighbor interpolation. In theory, this is equivalent to multiplying a matrix involving the upsampling to both sides of (1). This treatment results in very different behavior in the presence of noise. This is analogous to preconditioning techniques [18] for solving linear systems. Our model does not require that the LRPs fill the whole detector plane, for either regular or irregular layouts. We specifically model the gaps between physical pixels as *null* values, which better matches the information that is actually acquired by the sensor. Our novel error back projection algorithm iteratively recovers the super-resolved image for arbitrary pixel layouts, either regular or irregular.

2 OPTICS AND DETECTOR PROPERTIES

Super-resolution has been addressed mostly from a computational point of view, focusing on the conditioning of the reconstruction problem with respect to noise and quantization errors [7], [25]. Little attention is paid to the blur functions of real lenses and the sampling properties of real sensors. These are often approximated by a Gaussian blur for the lens and a box function for the sensor. In this section, we

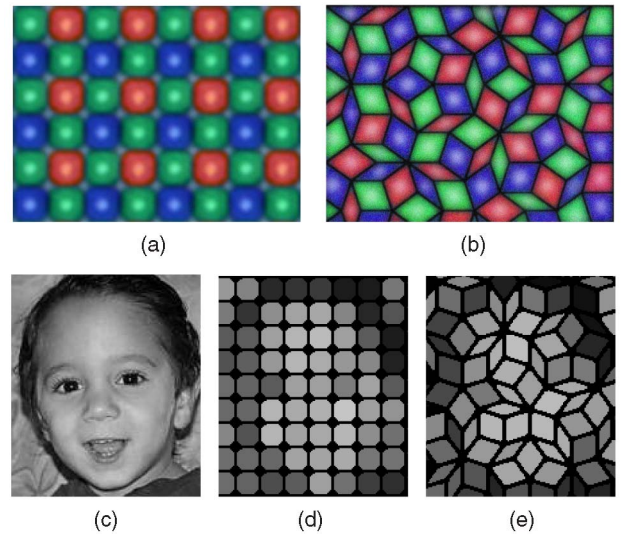


Fig. 1. A regular pixel layout and a Penrose pixel layout on the detector plane. (a) A microscopic view of a Sony 1/3'' sensor (part). Our method models the nonsquare shape of the pixels as well as the gaps between them. (b) A hypothetical aperiodic Penrose pixel layout. (c) An example image at the sensor surface (irradiance). (d) Spatial integration for the conventional layout. (e) Spatial integration for the Penrose layout.

discuss lens blur and spatial integration, and their impact on the performance of super-resolution of real images.

2.1 Modulation Transfer Function and Super-Resolution

The modulation transfer function $MTF(\psi)$ of a lens system describes how sharp a spatial sine wave will appear on the image plane prior to being sampled by the sensor. The MTF drops as the spatial frequency increases and limits the achievable magnification factor for super-resolution. To see this, consider a 1D case in which a perfect $[0, 1]$ vertical edge passes through a pixel. For a quantization level of ϵ and a pixel size of 1, the minimum displacement that produces a detectable change (above the quantization level) in the measured pixel intensity would be $d_{min} = \epsilon$. This implies that $\frac{1}{\epsilon}$ different values are obtained as the edge moves across the pixel. If the contrast level drops to $[0.25, 0.75]$, however, then the minimum detectable displacement would be $d_{min} = 2\epsilon$, thereby halving the number of possible different values.

For periodic tiling such as square pixel tiling, the maximum displacement before the complete image repeats is the tiling period length (one pixel for square pixel tiling). To collect more images, one must use shorter and shorter displacements. As the contrast level drops, however, small displacements will cause fewer and fewer pixels to change value (above noise level and quantization error). This is a fundamental limit for super-resolution using periodic tiling. By contrast, one can collect many measurements with aperiodic tilings using a series of half-pixel displacements in any given direction. Later, we will show that this leads to a much better conditioned system of equations for super-resolution.

2.2 Modeling Real Lens Blur

The MTFs of real lenses vary with distance from the optical center, orientation, focus, and of course aperture changes. Figs. 2a and 2b show the measured MTF of a Carl Zeiss




Fig. 2. MTF of the Carl Zeiss Planar T*1.4/50 standard lens. The plot shows the MTF value for 10 (top), 20, and 40 lp/mm at two different orientations as a function of distance from the optical center. Clearly, the MTF is not uniform and is neither shift invariant nor orientation invariant. Therefore, it is not well modeled by a Gaussian.

standard lens (Planar T*1, 4/50) for a specific focal distance, two different apertures and three spatial frequencies, as a function of the distance from the optical center. The solid and dashed lines refer to different orientations. The value of the MTF clearly varies for different image locations and orientations. Therefore, the blur of a real lens cannot be expected to be shift-invariant or Gaussian. Image deblurring assuming a Gaussian blur function results, in the best case, in a sharpening operator similar to the Photoshop unsharp mask. In other cases, it can lead to unwanted artifacts. For this reason, we do not attempt to account for lens blur. Instead, we only recover the super-resolved image at the sensor plane.


2.3 Modeling the Pixel Blur

Modeling the pixels of a sensor is an easier task than modeling the lens. Usually the pixel integration blur function is modeled as a 2D box function. This model assumes that pixels are uniform, perfectly square in shape, and that there are no gaps between pixels. Real sensors, however, are not perfectly square and abutting. When microlenses are used, the shape of the microlens, usually circular, determines the effective area (see Fig. 1a). When microlenses are not used, the shape of the pixels can vary significantly. Full frame and interline CCD sensors usually have square or rectangular pixels, but with different fill factors. The FujiFilm Super CCD has octagonal pixels and a diagonal layout. The shape of the active area of CMOS sensor pixels can vary from a small square to an “L” shape [31].

The shape of real pixels is therefore not modeled exactly by a regular grid of square pixels. In some cases, such as a sensor with low fill factor, the conventional modeling can be misleading; one example occurs when a strong but sharp feature falls in the gap between pixels. We want to model exactly what each pixel measures, and nothing more. Provided that we know the layout of the sensor, we can represent this knowledge in an image such as Fig. 1b. This is a high-resolution image, but it represents the information obtained by a lower resolution sensor as well as the layout of this sensor. In the next sections, we use this representation to model an aperiodic layout that cannot be approximated by a simple shift invariant kernel.

3 PENROSE TILING

Penrose tiling is an aperiodic tiling of the plane presented by R. Penrose in 1973 [19]. Figs. 1b and 3a show the



(a)

(b)

Fig. 3. Rhombus Penrose tiling is an aperiodic tiling consisting of two rhombi in five different orientations. This layout is neither regular nor amorphous as seen in its frequency domain image^c that clearly shows its five fold rotational symmetry.

rhombus Penrose tiling, which consists of two rhombi, each placed at one of the five different orientations according to specific rules [19]. The ratio of the number of thick to thin rhombi is the Golden Number $\frac{1+\sqrt{5}}{2}$, which is also the ratio of their areas. Unlike regular tiling, Penrose tiling has no translational symmetry—on an infinite plane, it never repeats itself exactly.¹ For the purpose of super-resolution, this means that it is theoretically possible to integrate and sample the infinite plane indefinitely without repeating the same pixel structure. In practice, this allows the capture of a significantly larger number of different images than is possible with a regular grid. Moreover, all images can be optimally displaced approximately half a pixel apart and still be different. By contrast, a regular square tiling forces the maximal delta between equally spaced displacements in x and y to be at most $\frac{1}{M}$, where M is the linear magnification factor. As mentioned in the previous section, this is very harmful to super-resolution. The rhombus Penrose tiling shown in Fig. 1b is a good candidate for a hardware color sensor implementation because it is 3-colorable [32] and has only two simple tiles. This is the primary reason that we selected this particular aperiodic tiling. By contrast, using random tessellation, such as a Voronoi tessellation over a random set of centers, results in numerous different tiles with different electric and photometric properties and will be very difficult to make.

4 OUR MODEL AND ALGORITHM FOR SR

This section describes our model and algorithm, which aims to obtain the best possible results for real photographic systems. Similarly to the Jitter-Camera [10], we assume the motion of the camera to be translational on a plane parallel to the detector plane. We also assume that the images are captured (or otherwise selected) in a *controlled* manner such that the displacements are equal in both the horizontal and the vertical directions and are exactly $\frac{1}{M}$ apart (for regular grid), where M is the linear magnification factor. The meaning of the magnification factor for irregular pixel layouts is somewhat ambiguous. We therefore refer to the resolution (and hence the magnification factor) as the number of pixels in a given area. Two different layouts are considered to have the same resolution if they have the

1. The Penrose tiling never repeats itself on an infinite plane, but any finite portion of the Penrose tiling can repeat infinitely many times [19].




Fig. 4. Upsampling and resampling. Upsampling is done by placing a regular high-resolution pixel (HRP) grid over the actual shape of the low resolution pixels (LRP), shown as white areas, then assigning the value of the LRP to each of the HRPs covering it. HRPs that (mostly) cover black areas (non-photo-sensitive areas) are assigned the value *null*. Downsampling is an inverse procedure that integrates the nonnull HRP values to form the value of its underlying LRP. Resampling is the composition of downsampling and upsampling.

same number of pixels within a given area (up to a rounding error).

The shape of LRPs can also be different from each other, and gaps between pixels are allowed. As in [8], [7], [25], we also assume that the pixels of the sensor have uniform photosensitivity, which implies that the contribution of an HRP to an LRP is proportional to its area inside the LRP and vice versa. These assumptions greatly simplify our model and implementation, but, as we show later in this section, they can easily be relaxed.

4.1 Upsampling and Resampling

In our approach, the LRPs in each LRI may not be aligned on a regular grid. Nonetheless, we can still index each LRP in an LRI because we have full knowledge of the pixel layout. As soon as an LRI is captured, we immediately upsample it to the high-resolution grid to create an intermediate HRI. This intermediate HRI, not the original LRI, is involved in the computations that follow. As shown in Fig. 4, upsampling is done by placing a regular high-resolution pixel grid over the actual shape of the low-resolution pixels and then mapping HRPs to LRPs. HRPs that are not associated with any LRP are assigned the value *null* to differentiate them from the value zero. The assumption of pixel uniformity can be relaxed at this stage by multiplying the intermediate HRI with a weight mask to compensate for any intrapixel nonuniformities. For fronto-parallel translational motion with image displacements equal to $\frac{1}{M}$, it turns out that the registration of the intermediate HRI is simply an integer shift of the origin. If the motion assumptions do not hold, an additional warping step is required after the upsampling. We denote the upsampling operator by $\uparrow_{T_i, G}$, where T_i is the transformation for registration and G is the sensor layout map.

Our algorithm also includes an error back projection procedure. It requires a resampling operator (Fig. 4), denoted by $\downarrow_{T_i, G}$, which simulates the image formation process to produce new intermediate HRIs given an estimate of the super-resolved image. The resampling operator can be viewed as a downsampling operator followed by an upsampling operator. An alternative way to view the upsampling and resampling operators is to view the downsampling operator as filling each LRP with the average value of the HRPs inside it, and the upsampling/resampling operator as filling the HRPs inside the same LRP with their average value. In practice, the computation is done “in-place” and no actual downsizing takes place. The

resulting images are hypotheses of the intermediate HRIs, assuming that the super-resolved image is the correct one.

4.2 Error Back Projection Algorithm

Our super-resolution algorithm is a variant of the well-known error back projection super-resolution algorithm [23]. Unlike the traditional algorithm, which downsamples the images into a low resolution array, our algorithm is performed entirely on the high-resolution grid. Using the concepts in the previous section, we summarize our algorithm as follows:

Algorithm 1

Inputs:

- L_1, \dots, L_{M^2} : Low resolution images (with roughly N^2 pixels).
- T_1, \dots, T_{M^2} : Transformations for registering the LRIs.
- $M \in \mathcal{N}$: Magnification factor.
- G : Sensor layout map.

Output:

- S : Super-Resolved image ($NM \times NM$).

Processing:

- 1) Upsample: $H_i = L_i \uparrow_{T_i, G}$, $i \in [1, \dots, M^2]$.
- 2) Initialize: $S^0 = \frac{1}{M^2} \sum_{i=1}^{M^2} H_i$.
- 3) Iterate until convergence:
 - a. $S^{k+1} = S^k + \frac{1}{M^2} \sum_{i=1}^{M^2} (H_i - S^k \downarrow_{T_i, G}) \uparrow_{T_i, G}$.
 - b. Limit: $0 \leq S^k(x, y) \leq \text{MaxVal}$.

Note that *null* elements are ignored when computing the average values. Step 3)b represents the prior knowledge about a physically plausible image, where MaxVal is determined by the optical blur and the A/D unit. The difference between our algorithm and the conventional back projection algorithm (with a *rect* kernel) lies in the upsample stage. Our upsampling operator $\uparrow_{T_i, G}$ preserves sharp edges between pixels at the high-resolution grid, whereas the conventional algorithm applies the blur kernel globally. If warping is required, it is performed on the intermediate HRI after the upsampling.

5 ANALYSIS

In this section, we analyze our super-resolution algorithm (Algorithm 1) and the irregular Penrose pixel layout. We begin by analyzing the error bounds for our algorithm and showing that it converges more closely to the ground truth for an appropriately chosen aperiodic pixel layout than for a periodic one. We also discuss the condition number of the RBA system of equations and show that the Penrose pixel tiling leads to a better conditioned system, which is thus more robust to noise. Finally, we look at the information content of Penrose Pixel and regular pixel layout LRIs and present experimental results suggesting that sets of Penrose Pixel LRIs contain more information.

5.1 Convergence of the Algorithm

For the j th LRP of the i th LRI, its downsampling and upsampling operators can be represented by $N_j^{-1}(p_j^i)^T$ and




Fig. 5. Comparison of (a) the condition number and (b) the numerical rank for Penrose pixel and square pixel layouts. Fixed threshold (0.05) and different magnification factors. (c) Numerical rank for the Penrose pixel and square pixel layouts. Fixed magnification factor ($\times 8$) and different thresholds. The Penrose layout is more numerically stable with respect to both the condition number and the numerical rank.

p_j^i , respectively, where N_j is the number of HRP inside the j th LRP and p_j^i is a binary vector: $p_j^i(k) = 1$ indicates that the k th HRP is inside the j th LRP of the i th LRI, and $p_j^i(k) = 0$ if not. The superscript T denotes transpose. Then the resampling operator, as the composition of downsampling and upsampling, associated with the j th LRP of the i th LRI is $N_j^{-1}p_j^i(p_j^i)^T$. Hence, the resampling matrix representing the resampling operator $\downarrow_{T_i, G}$ in Algorithm 1 can be written as

$$R_i = \sum_j N_j^{-1} p_j^i (p_j^i)^T. \quad (2)$$

So, every intermediate HRI H_i aligned to the high-resolution grid is connected to the ground truth image S via the following linear system:

$$H_i = R_i \cdot S + n_i, \quad i = 1, 2, \dots, M^2, \quad (3)$$

where n_i are the noise from L_i . As the pixel shape and layout of LRPs is irregular, it is difficult to write down exactly how R_i looks. Also, as mentioned in the previous section, in practice we do not explicitly compute the down/up resampling matrices.

The iteration 3)a in Algorithm 1 can be written as

$$S^{k+1} = S^k + \frac{1}{M^2} \sum_{i=1}^{M^2} (R_i \cdot S + n_i - R_i \cdot S^k), \quad (4)$$

which can be rewritten as

$$S^{k+1} - S = (I - \bar{R})(S^k - S) + \bar{n}, \quad (5)$$

where $\bar{R} = \frac{1}{M^2} \sum_{i=1}^{M^2} R_i$ and $\bar{n} = \frac{1}{M^2} \sum_{i=1}^{M^2} n_i$. So,

$$S^k - S = (I - \bar{R})^k (S^0 - S) + \left[\sum_{j=0}^{k-1} (I - \bar{R})^j \right] \bar{n}. \quad (6)$$

We can prove that the spectral radius of $I - \bar{R}$ is usually less than one (see the Appendix), then $\lim_{k \rightarrow \infty} (I - \bar{R})^k = 0$ and \bar{R} is nonsingular with $\bar{R}^{-1} = \sum_{j=0}^{\infty} (I - \bar{R})^j$. Then, from (6), we have that

$$\lim_{k \rightarrow \infty} S^k = S + \bar{R}^{-1} \bar{n}, \quad (7)$$

which means that iteration 3)a in Algorithm 1 converges to an HRI which deviates from the ground truth by $\bar{R}^{-1} \bar{n}$.

5.2 Error Analysis and Numerical Stability

From (7), we may expect that the iterations result in a super-resolved image which deviates from the ground truth by $\bar{R}^{-1} \bar{n}$. Note that \bar{n} can be viewed as the empirical estimation of the mean of the noise. Therefore, when the noise in the LRIs is of zero mean (and so is n_i , as there is a linear transform between them), we can expect that a high-fidelity super-resolved image is computed. If we choose an appropriate pixel layout so that the norm of \bar{R}^{-1} is small, then the deviation can be effectively controlled regardless of the mean of the noise (note that $\|\lim_{k \rightarrow \infty} S^k - S\| = \|\bar{R}^{-1} \bar{n}\| \leq \|\bar{R}^{-1}\| \|\bar{n}\|$). As $\|\bar{R}^{-1}\|$ is large when \bar{R} is close to singular, we should choose an appropriate detector pixel layout such that \bar{R} is far from singular.

According to the above analysis, we should choose pixel layouts that result in more linearly independent equations in the system (3). The traditional regular tiling repeats itself after a translation of one LRP (two LRPs if we account for the Bayer pattern in color sensors). Lin and Shum [25] also showed that if five LRPs cover the same set of HRPs, then their equation set must be linearly dependent. Thus, using regular (and square) tilings usually results in an insufficient number of independent equations. To overcome this difficulty, we try to change the regular tiling to other kinds of tilings. An intuition is to use aperiodic tilings.

In an attempt to quantify the difference between the Penrose LRIs and the regular LRIs, we empirically computed the condition number as well as the numerical rank (the number of singular values larger than a threshold) of \mathbf{P} from (1) for different magnification factors. Figs. 5a, 5b, and 5c show the condition numbers and the numerical ranks with respect to different thresholds for Penrose pixels and periodic square pixels as a function of the magnification factor. We can see for the same magnification factor, the condition numbers for the Penrose layout are up to an order of magnitude lower than those for the square pixels, and the numerical ranks for the Penrose layout are larger than those for the square pixels. We can also see that the condition number for the Penrose layout is much more stable than

that for the square pixels in that it does not suffer from numerical instability at integer magnification factors [7], [25]. The numerical rank is similar. These plots show that the linear system for the Penrose layout is much more stable than that for the regular square layout.

5.3 Improved Conditioning Using the Resampling Operator

The previous section presented empirical evidence of the superior conditioning of Penrose Pixel RBA. We now explain why the resampling operator can result in better SR performance than traditional reconstruction-based methods. Equation (3) leads to the following overdetermined linear system:

$$H = R \cdot S + n, \quad (8)$$

where $H = (H_1^T, H_2^T, \dots, H_{M^2}^T)^T$, $R = (R_1^T, R_2^T, \dots, R_{M^2}^T)^T$, and $n = (n_1^T, n_2^T, \dots, n_{M^2}^T)^T$. By contrast, the traditional formulation for reconstruction-based SR is

$$L_i = D_i \cdot S + e_i, \quad i = 1, 2, \dots, M^2. \quad (9)$$

Here, e_i is the noise from L_i and the matrix D_i is the representation of the downsampling operator in Fig. 4, leading to the following overdetermined linear system:

$$L = D \cdot S + e, \quad (10)$$

where $L = (L_1^T, L_2^T, \dots, L_{M^2}^T)^T$, $D = (D_1^T, D_2^T, \dots, D_{M^2}^T)^T$, and $e = (e_1^T, e_2^T, \dots, e_{M^2}^T)^T$. The major difference between (3) and (9) is that D_i is roughly a submatrix of R_i (R_i has many more rows than D_i). So, D is also roughly a submatrix of R . It is well known in matrix theory [22] that, in this case, the minimum singular value of R , defined by $\sigma_{\min}(R) = \min_{x \neq 0} \|Rx\|/\|x\|$, will be no smaller than that of D . This increase of the minimum singular value can greatly reduce the condition number of the corresponding linear system, particularly when the original system (10) is ill-conditioned. Although the maximum singular value of R is also larger than that of D , our numerical simulation shows that with very high probability, the condition number, defined as the ratio of the maximum singular value to the minimum one, is reduced by adding rows to a matrix. When the original matrix is very well conditioned, it happens occasionally that the condition number increases *slightly*, so the robustness of the system is virtually unchanged. Generally speaking, using the resampling operator instead of the traditional downsampling operator can result in a much better conditioned system. As a consequence, better SR results are possible.

5.4 Information Content for Penrose Pixel LRI's

Intuition suggests that because Penrose pixel views can all be displaced by half LRP intervals, the differences between adjacent views would be larger than for regular pixels, which must be displaced by $\frac{1}{M}$ ($M > 2$). When subject to quantization error, this should provide an advantage to the Penrose pixels over the regular one. How can we verify this quantitatively? Here, we present a rough empirical estimate of the relative information content of regular and Penrose pixel layouts. We do this using gzip to compress images captured using both layouts. Gzip uses LZ77, which for large file sizes converges to optimal compression.




Fig. 6. File sizes of sample (low resolution) image sets for Penrose and regular tiling compressed by gzip. The compressed size of the Penrose tiling set is consistently larger than the size of the compressed regular set. The compressed size of $\times 8$ (64 images) and $\times 16$ (256 images) for the same quantization level is very similar. This indicates (but does not prove) that the Penrose tiling possesses more information than the regular tiling for the same number of pixels.

We assume that, for real images (in contrast to special cases like pseudo-random data), lossless compression is monotonic, i.e., a larger compressed file implies larger data complexity. We assume neither optimality nor linearity of the complexity with respect to the file size. We also ignore the layout description of both the Penrose and regular pixel tiling because Penrose and regular tillings of the infinite plane can be described by a finite (and relatively small) set of rules. Because the layouts do not match, it is not clear in which order to serialize the pixels, as this may introduce a bias. We solve this problem by using the nonoptimal, but nonbiased, LRI representation in the high-resolution grid and simply compressing the views in the high-resolution grid.

For our comparison we used the regular and Penrose layouts shown in Figs. 1d and 1e. Both are nonsquare pixels, both have gaps between pixels, and both have the same average area. We did the following:

1. Create 64 images for $\times 8$ magnification factor for Penrose and regular pixel layouts. The displacement step is $\frac{1}{8}$ pixel for the regular tiling and $\frac{1}{2}$ pixel for the Penrose tiling.
2. Concatenate all views into a single file to allow the compression algorithm to apply dictionary entries obtained in one view to other views.
3. Compress the concatenated file using gzip and record the absolute file size.
4. Repeat steps 1-3 for quantization levels of 8, 6, 4, and 2 bits.
5. Repeat steps 1-4 for magnification factor of 16x, and 256 images.

Fig. 6 displays the compressed file size (in K-bytes) for different quantization levels and magnification factors. The first thing to notice is that although the number of images for the magnification factor of 16x was four times larger than those for 8x magnification, the *compressed* file size is very similar, even slightly smaller. This indicates that the amount of information in these sets did not change significantly (slight variations are expected due to convergence properties of the algorithm and technical reasons

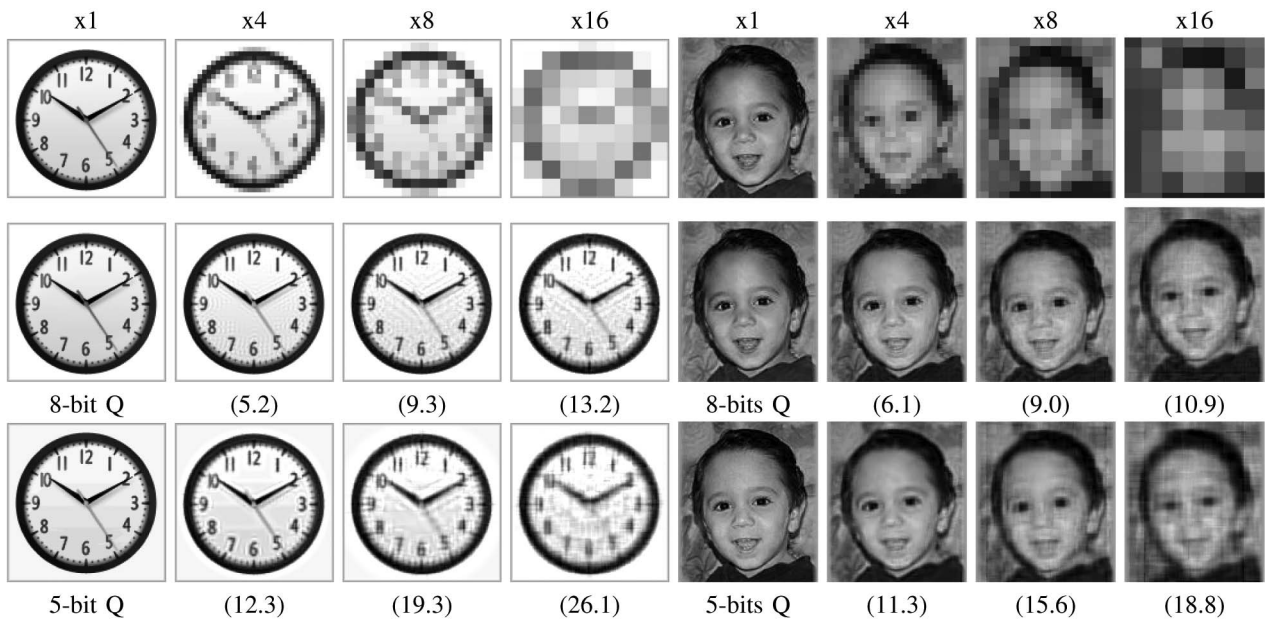


Fig. 7. Effects of Quantization Error: Super-resolution results for the “clock” and “face” images using regular tiling. Top left corner: Original image. Top row: LRIs with different magnification factors (scaled). Center and bottom rows: Super-resolution results for quantization levels of 8 and 5 bits, respectively. The results *gradually* degrade as the quantization error and magnification increase. Parentheses are the RMS errors. M^2 input images were used.

such as the dictionary size). We also see that the Penrose tiling LRIs are consistently and quite significantly compressed to a larger file size than the regular tiling LRIs, and that the difference grows with the quantization level. This suggests that the Penrose tiling LRIs do contain more information than the regular ones. This is due to the combined effect of the displacement size, the pixel aperiodicity, and the existence of two different pixel shapes.

6 TESTING AND EVALUATION

We evaluated our approach with simulations and real image tests. For our first experiment, we simulated the entire super-resolution process for square and Penrose pixels. The integration of each pixel was approximated using the sum of the pixels in the high-resolution grid that are enclosed within the low-resolution grid pixel area. As we do not have an actual Penrose pixel sensor, our second experiment strives to be as close to real world conditions as possible. We first captured a sequence of high-resolution real images (each with its own unique noise values) and then integrated pixel values to simulate a Penrose image. The last experiment is a start-to-finish real image super-resolution test.

To fully utilize the advantage of the aperiodic layout, and to overcome noise, we usually used more images than unknowns. The advantage of using an overdetermined system is shown in Fig. 12. For the case of square pixels and quantization error only, we used M^2 input images, where M is the magnification factor. This is the maximum number of *different* images we can obtain using a displacement of $\frac{1}{M}$. The number of input images used in each test appears in the captions of the relevant figures.

6.1 Regular Pixels Quantization and Noise Tests

In our first simulation, we applied our algorithm to LRIs synthesized from ground truth HRIs of a clock and a face.

We used regular grids with linear magnification factors of 1 to 16, and quantization levels of 8 and 5 bits. No additional noise was added. Fig. 7 shows our super-resolution results and RMS errors (compared to the original image). Though there is a gradual degradation with increasing magnification and quantization error, the super-resolution algorithm performs very well. This matches our analysis for zero mean (quantization) noise.

Next, we added Poisson noise (which better models real noise) to the input images. Fig. 8 shows the super-resolution result for the “face” image using additive Poisson noise with mean = 5 and 10 gray levels, followed by 8-bit quantization. Unlike the zero-mean quantization error, the nonzero mean Poisson noise significantly degrades the quality of the results. The results can be improved by using many more images than the theoretical minimum requirement, as shown in the bottom row of Fig. 8.

6.2 Penrose Pixels Quantization and Noise Tests

We repeated the last two tests for two Penrose tiling pixel layouts. The magnification factors were roughly equivalent to 8 and 16, and the quantization level was 8-bit. Unlike for regular pixels, we used displacements of approximately 0.5 pixels and were able to use more images than was possible with the regular grid. The results shown in Fig. 9 are clearly better than the results obtained with the regular grid, shown in Figs. 7 and 8.

To better quantify the results, we used a concentric test target having variable spatial frequency contrast.² We added low-level noise to each image to create quantization variations. Then, we applied our algorithm and the conventional back projection algorithm under *exactly the same conditions* and using the same number of input images. Fig. 10 shows that our algorithm improves the linear magnification by

2. The contrast of real lenses decreases as the spatial frequency increases.




Fig. 8. Noise evaluation test using regular tiling. The top two rows show the results of super-resolution with different magnification factors and Poisson noise of mean 5 (top) and 10 (middle) gray levels. M^2 LRIs were used and all computations were run for several thousand iterations or until convergence. The amplification of noise is quite clear and the results are very different from those with zero-mean quantization error shown in Fig. 7. The bottom row shows the results of super-resolution with magnification of $\times 8$ and noise mean of 10 running for 1,000 iterations, with different numbers of input images (256 to 1,600). The results are much better than that obtained by using the minimum 64 input images (the boxed image).

roughly a factor of two (for the same RMS errors) compared to the conventional back projection algorithm with regular pixels, and by over a factor of four when Penrose pixels are also used. In this test, we also compared to the RMS errors of the conventional algorithm with a noninteger magnification factor to rule out the possibility that the difference is due to the integer magnification used in our algorithm [7]. As shown in Fig. 10 the integer magnification factor did not affect the results for the back projection algorithm.

Fig. 12 compares the RMS error as a function of the number of images for regular and Penrose tiling, respectively. The magnification factor was eight and the same algorithm (Algorithm 1) was applied to both layouts. While the regular layout improved slightly when overconstrained, the Penrose layout improved by over four times. It is interesting to see that the regular layout was actually better when the system was severely underconstrained.

Fig. 13 shows different image types and the convergence of the algorithm (RMS error) as a function of the number of iterations.

For our last simulation example, we compared our algorithm to an externally obtained result of [21], [7] using an image from the FERET database [6]. In Fig. 11, the improvement from our approach is clearly visible.

6.3 Real Images with Simulated Binning Test

In this test, we captured 576 real images with a Nikon D70 camera on a tripod. We computed the LRIs by integrating each image with the map G , then quantizing the result. Thus, the resulting LRIs had unique noise due to sensor noise, quantization, and JPEG compression. This process is

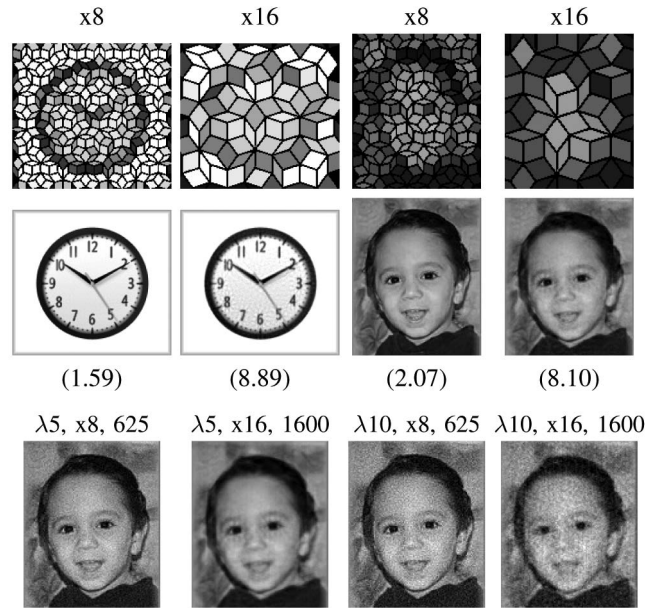


Fig. 9. Penrose Pixel super-resolution results. (Top) Input images for magnification factors of 8 and 16. (Middle) 8-bit quantization result. Compare to the corresponding images in Fig. 7. (Bottom) 8-bit quantization, and Poisson noise results. Compare the image labeled “ $\lambda 10 \times 8 625$ ” to its corresponding image in Fig. 8. The *overconstraint* number of *different* LRIs used was 625 for $\times 8$ and 1,600 for $\times 16$ in all cases.

very similar (though noisier) to pixel binning done at the analog level. As with pixel binning in real sensors, large simulated pixels have lower noise than small integrated pixels do. The LRIs were also subject to slight misalignment due to shake by the flipping mirror in the camera. Capturing more images than the required minimum reduces the effect of slight misalignments. Fig. 14 shows the results of applying our super-resolution algorithm to the LRIs for regular and Penrose layouts. The advantage of the Penrose layout is clear.

6.4 Real Scenario Test

For our real-world test, we captured a sequence of images using a B/W version of the Sony $1/3''$ sensor shown in Fig. 1a. Using the lens resolution and pixel size and shape, we created a sensor model for $\times 5$ magnification (which is above the nominal lens resolution). We model square pixels with trimmed (null) corners to match the actual pixel shape, including the microlens. We then moved a test image, in a controlled manner, in front of the camera and captured 5×5 input images at 25 different displacements. To reduce noise, we averaged 20 frames for each input image.³ Fig. 15a shows one of the input images and a magnified insert. Fig. 15b shows an actual intermediate HRI. The black dots are the null values at the corners of each pixel. Fig. 15c shows the super-resolution result. Note that even fine details such as the dots above the letter “i” and in the exclamation marks were resolved.

3. This is possible in a controlled environment such as the “Jitter Camera” [10], and saves a lot of storage space and computation time. In uncontrolled environments, all captured images should be fed directly into the super-resolution algorithm to reduce noise and misalignments artifacts.

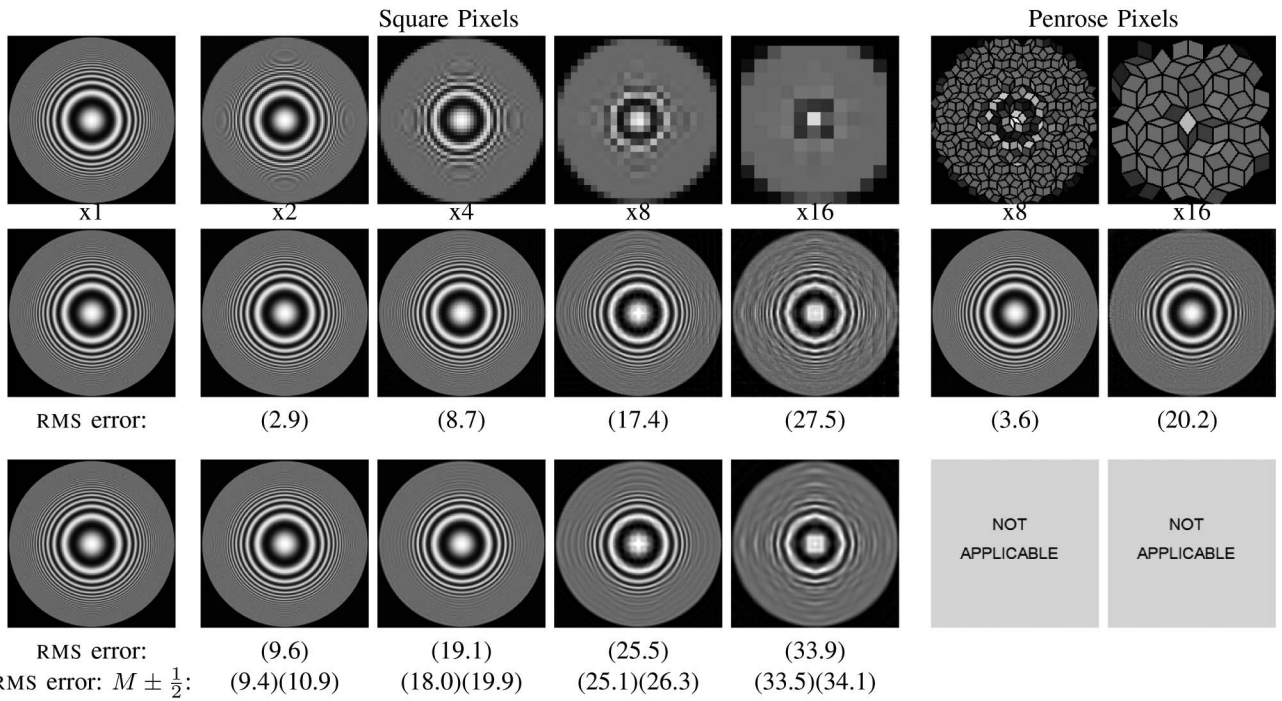


Fig. 10. Test target comparison. Top: Input images for regular and Penrose pixel layouts, with magnification factors of 8 and 16, respectively. Middle: Super-resolution results using our back projection algorithm for the regular and Penrose pixel layouts. Bottom: Super-resolution results using the conventional back projection algorithm for the regular layout (with matched Gaussian kernel). Below are the RMS errors for noninteger magnification factors of $M \pm \frac{1}{2}$. In all cases, the number of LRIs used was 1,600.

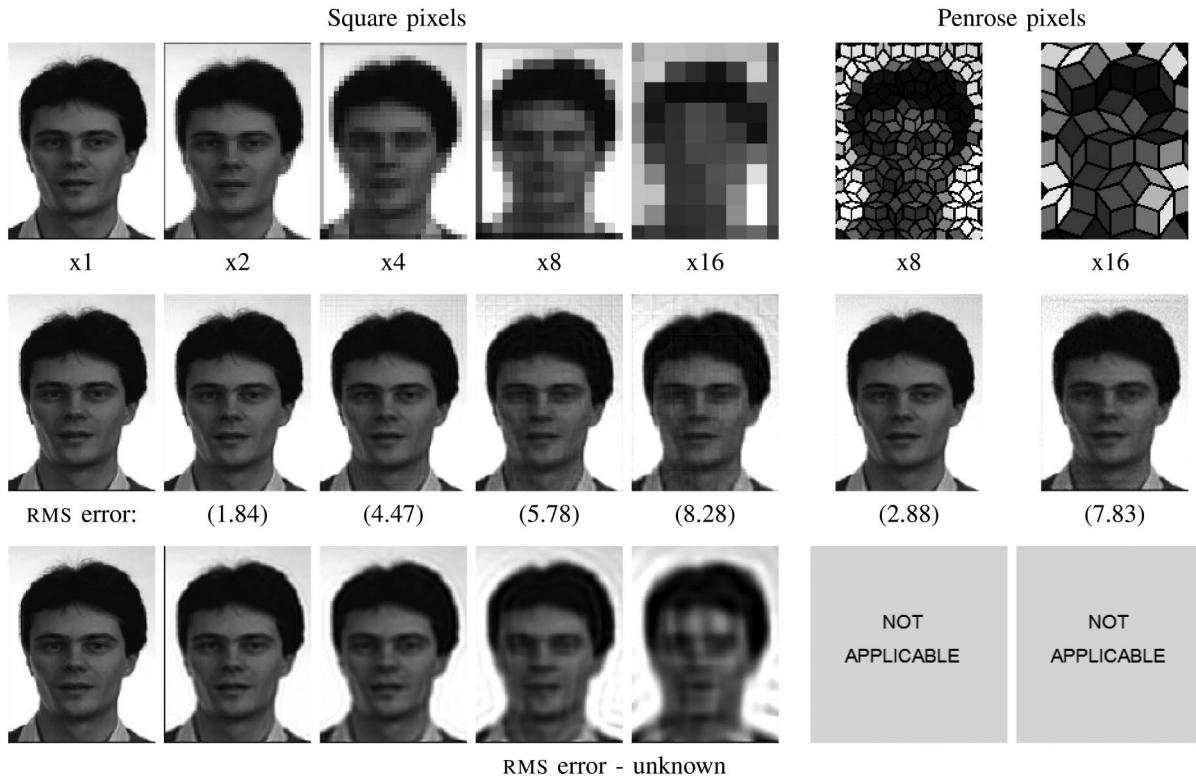


Fig. 11. Comparison to external result (FERET DB image). Top: Low resolution images at different magnification factors. Middle: Our results for square and Penrose pixel layouts. Bottom: Result of super-resolution using [21] (image taken from [7]). For the square pixels, M^2 LRIs were used. For the Penrose Pixels, overconstraint sets of 256 and 625 different images (due to the aperiodic tiling) were used for magnification of $\times 8$ and $\times 16$, respectively.

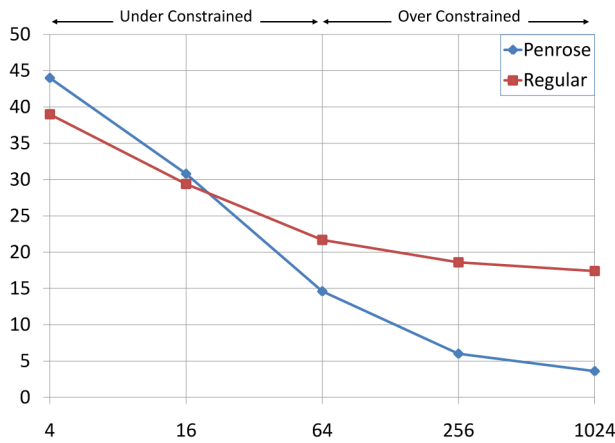


Fig. 12. RMS error versus number of images for $\times 8$ magnification factor. (a) Regular layout with square pixels. (b) Penrose layout. The Penrose layout clearly better utilizes the additional images.

7 DISCUSSION

So far we have only addressed the super-resolution related aspects of Penrose tiling. We have mentioned before that Penrose rhombus tiling is 3-colorable, allowing the use of RGB color filter arrays on the sensor. Which coloring to use

and the best way to demosaic the image are open problems. An interesting aspect of Penrose Pixels is their irregular sampling. The acquired images are not subject to strong moiré effects that can plague conventional digital photography, particularly in video. Also, Penrose rhombus tiling is only one possible aperiodic tiling, which we selected mainly for its simplicity. Further research is needed to determine which tiling, if any, performs best.

Before we conclude our paper, we briefly address the plausibility of a hardware Penrose pixel implementation. At first glance, manufacturing an image sensor that uses an aperiodic pixel layout might seem implausible. In today’s sensor technologies (CMOS and CCD chips), control signals and power supplies are routed to each pixel using metal wires. These wires are opaque and typically run on top of the silicon substrate containing the photodetectors in each pixel. On a regular grid, wires can be run between pixels to minimize their negative impact on the pixels’ light gathering efficiency. This is not true for Penrose tiling.

Penrose pixel routing becomes much simpler if we assume a back-illuminated CMOS sensor. In such devices, the chip is thinned and mounted upside down in the camera so light enters from the back of the chip. The metal layers are now underneath the photodetectors, so they do

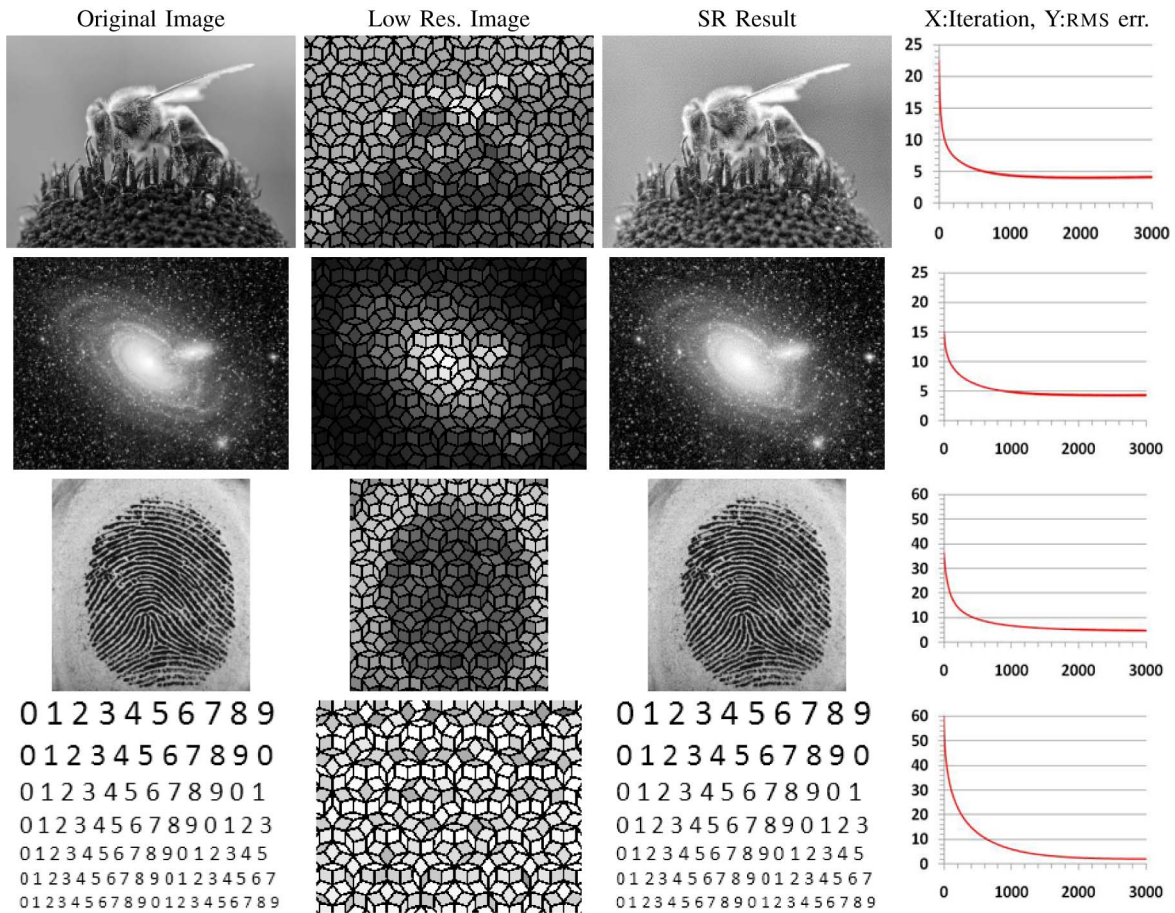


Fig. 13. Super-resolution results and convergence plots for different image types. We can see that the algorithm converges very fast during the first few iterations and then the convergence slows down (this is typical of error back projection algorithms). However, the initial error depends on the image content and contrast and affects the rate of convergence. All images were subject to quantization error and low additive Poisson noise. Magnification factor for all images is $\times 8$, and the number of LRIs was 256. Image source: Bee image—Wikimedia, Image by Jon Sullivan. Galaxy image—Hubble heritage gallery. Finger print image—FVC2000 db, University of Bologna.

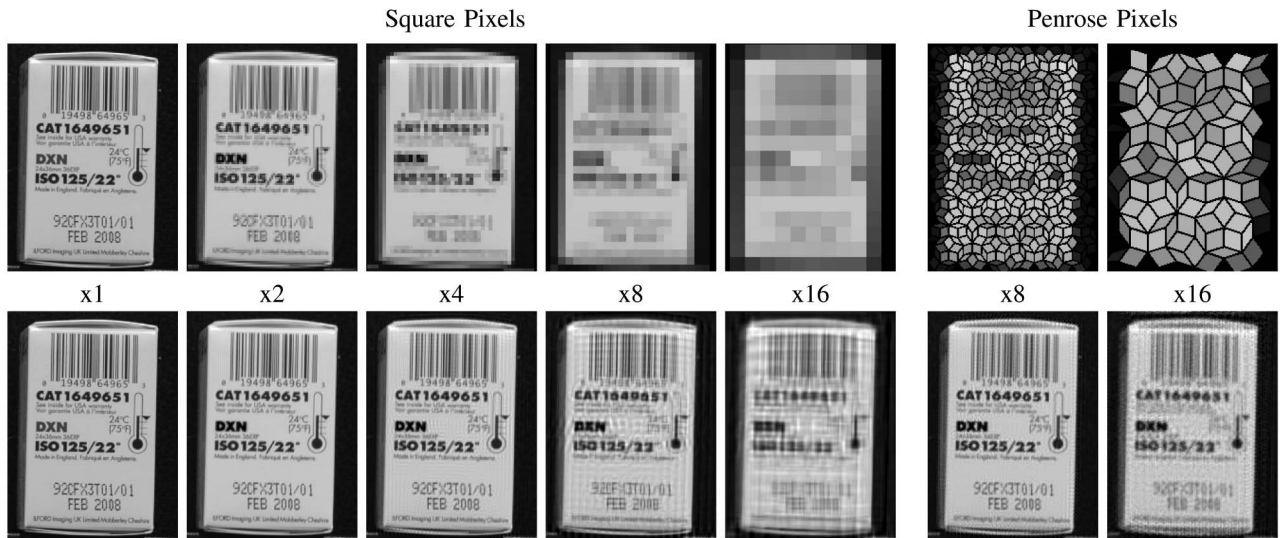


Fig. 14. Our algorithm applied to real images (each with its own noise) with simulated pixel integration. Top: One of 576 input images for the regular and Penrose pixel layout with magnifications of 1 to 16. Bottom: Results of our super-resolution algorithm applied to both regular and Penrose pixel layouts. In all cases, all 576 LRIs were used.

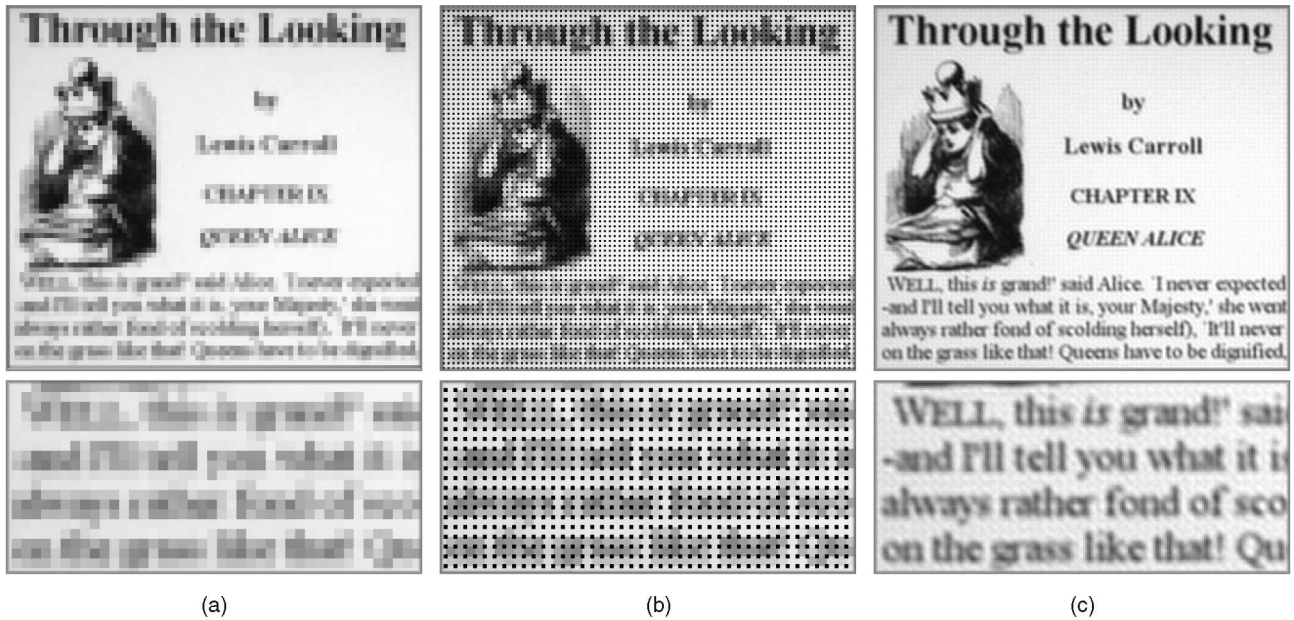


Fig. 15. Real image super-resolution results. (a) A real image captured by the camera (enlarged). (b) The input view for the super-resolution algorithm using our CCD model. The little black dots are gaps between pixels and have null value. (c) Super-resolution result. Notice the details in the magnified view, in particular the dots above the letter “i” and in the exclamation marks. A total of 500 LRIs were used at 25 displacements (20 images were averaged at each displacement).

not block light. In the past, manufacturing backside illumination (BSI) was relatively difficult and the technology was used only for very high-performance applications like astronomy. As pixel physical dimensions shrink, however, BSI is becoming attractive for consumer devices. Omnivision [2], one of the world’s leading manufacturers of CMOS image sensors, introduced an 8MP backside-illuminated sensor in 2008, and STMicroelectronics and Soitec announced a partnership in May 2009 to develop backside-illuminated image sensors for consumer products [3].

With no concerns about occluding pixels, Penrose pixel routing becomes much simpler. In a conventional CMOS image sensor, each row of pixels shares a “word line” and

each column shares a “bit line.” When the word line is asserted, the corresponding row of pixels drives their data (a voltage) onto their bit lines to be read out. Power supplies and other control signals can run in parallel to the word lines and bit lines. The wiring challenge for Penrose pixels is to associate each pixel with a unique wordline/bitline combination. Such a routing would be difficult for a frontside-illuminated Penrose Pixel sensor layout, but, using backside illumination, it becomes simpler.

Fig. 16 shows a simplified routing scheme for a Penrose Pixel image sensor. In this diagram, the Penrose pixels, drawn in black, are on the top of the sensor. The green and red lines, representing word and bit lines, respectively, are




Fig. 16. A Penrose pixel routing scheme. In this diagram, the Penrose pixels, drawn in black, are on the top of the sensor. The word and bit lines, represented by green and red lines, respectively, are on the underside. The gray circle on each pixel represents the connection point for signal wires. Each pixel must be connected to a unique word/bit line pair. As the diagram shows, this is possible with a sufficiently dense regular layout of word and bit lines, although some word/bit line pairs may have no associated Penrose pixel. Power supplies and other signals (not shown for clarity) would run parallel to the word or bit lines.

on the underside. The gray circle on each pixel represents the connection point for signal wires. Each pixel must be connected to a unique word/bit line pair. As the diagram shows, this is possible with a sufficiently dense regular layout of word and bit lines, although some word/bit line pairs may have no associated Penrose pixel.

Using the properties of this Penrose tiling, we can estimate the required density of routing wires. The ratio of the number of thick to the number of thin tiles in an infinite tiling is the Golden Ratio, $(1 + \sqrt{5})/2$. The acute angles in the rhombi are 36 degrees for the thin tiles and 72 degrees for the thick ones. For Penrose tiles of unit edge length, this implies a density of slightly over 1.23 tiles per unit area. Thus, the horizontal or vertical pitch of the word and bit lines must be at least $1/\sqrt{1.23} \approx 0.901$. Because the frequency of tiles varies locally, in practice we use a slightly higher density to ensure that all pixels can be routed. The bit line and word line pitch in Fig. 16 is 0.77. All pixels are connected to different word/bit line pairs, although some pairs are left unconnected.

Each of the two pixel shapes occur in five different orientations, so a maximum of only 10 unique pixel designs would be necessary. The finite number of neighboring pixel pair layouts could all be checked to prevent integrated circuit manufacturing design rule violations. Assuming that we place the pixels with custom software, standard IC wire routing tools could easily connect each pixel to the necessary wires (e.g., power supplies, a unique wordline/bitline combination, and so on) while ensuring other desirable properties like small signal wire lengths.

One might ask if it is feasible to fabricate an image sensor with two different diamond-shaped pixels. The irregular size of the photodetector itself is not a problem. Fujifilm, for example, has produced an image sensor with two oblong, differently sized photodiodes under a single microlens in each pixel [5]. We also require microlenses with shapes that match the diamond-shaped pixels. Such microlens arrays

can be produced using melting photoresist [13] in a similar way to hexagonal microlens array production [24].

Given the existing proven technologies described above, we are optimistic that it is possible to create an image sensor with irregularly shaped pixels and aperiodic tiling. Creating such an unconventional sensor will certainly involve some challenges. For example, current sensor designs rely on the similarity of periodic structures to reduce fixed pattern and random noise. A Penrose Pixel sensor will most likely exhibit more noise. Fortunately, for a high-end camera one can afford to use more sophisticated methods to overcome the noise, such as measuring the fixed pattern noise per pixel. For high-performance applications requiring high resolution and large pixels (for high sensitivity and dynamic range), we believe the benefits of Penrose Pixels will justify exploring a silicon implementation.

8 CONCLUSION

We present a novel approach to super-resolution based on an aperiodic Penrose tiling and a novel back projection super-resolution algorithm. Our tests show that our approach significantly enhances the capability of reconstruction-based super-resolution, as well as bringing it closer to bridging the gap between the optical resolution limit and the sensor resolution limit. We also argue that constructing a real Penrose tiling sensor is feasible with current technology. This could prove very beneficial for demanding imaging applications such as microscopy and astronomy. Another exciting possibility is to adapt current image stabilization jitter mechanisms [4] for use with super-resolution. Even a modest $4\times$ linear magnification would turn an 8MP camera into a 128MP one for stationary and possibly moving [10] scenes, without changing the field of view.

APPENDIX

In this appendix, we prove that the spectral radius of $I - \bar{R}$ is usually less than 1.

As R_i can be written as (2), R_i is symmetric and positive semidefinite. Moreover, the sums of the rows of R_i are either 0 or 1. \bar{R} , as the mean of R_i s, is then also symmetric and positive semidefinite, and the sums of its rows never exceed 1, i.e.,

$$\sum_q \bar{R}(p, q) \leq 1, \quad \forall p. \quad (11)$$

Then, by the Geršgorin disk theorem [22] and the nonnegativity of \bar{R} , the eigenvalues of \bar{R} lie in the union of the following disks:

$$D_p = \left\{ \lambda \mid |\lambda - \bar{R}(p, p)| \leq \sum_{q \neq p} \bar{R}(p, q) \right\}.$$

Note that disk D_p is inside

$$\tilde{D}_p = \left\{ \lambda \mid |\lambda| \leq \sum_q \bar{R}(p, q) \right\}.$$

So, we can see that the eigenvalues $\lambda(\bar{R})$ of \bar{R} satisfy $|\lambda(\bar{R})| \leq 1$ due to (11). Since \bar{R} is positive semidefinite, we actually have $0 \leq \lambda(\bar{R}) \leq 1$.

If 0 is an eigenvalue of \bar{R} , then there exists a nonzero vector v such that $\bar{R}v = 0$. Then, $v^T \bar{R}v = 0$, i.e.,

$$\sum_i \sum_j N_j^{-1} [(p_j^i)^T v]^2 = 0.$$

So,

$$(p_j^i)^T v = 0, \quad \forall i, j.$$


This means that if the HRP values are chosen as those of v , then for every displacement of the LRIs, the resampled HRI is always a zero image. This is quite impossible thanks to the irregularity of the shape and layout of the LRP. So, we actually have $0 < \lambda(\bar{R}) \leq 1$. Then, we conclude that the eigenvalues $\lambda(I - \bar{R})$ of $I - \bar{R}$, which is $1 - \lambda(\bar{R})$, satisfies $0 \leq \lambda(I - \bar{R}) < 1$.

ACKNOWLEDGMENTS

The Penrose tiling postscript program is by courtesy of Björn Samuelsson. This work was done while Bennett Wilburn was at Microsoft Research Asia.

REFERENCES

- [1] Hexagonal Image Processing Survey, <http://www-staff.it.uts.edu.au/wuq/links/HexagonLiterature.html>, 2011.
- [2] Omnivision Technologies, <http://www.ovt.com>, 2011.
- [3] STMicroelectronics and Soitec Join Forces to Develop Next-Generation Technology for CMOS Image Sensors, <http://www.st.com/stonline/stappl/cms/press/news/year2009/t2379.htm>, 2009.
- [4] www.dpreview.com/reviews/minoltadimagea1/, 2011.
- [5] www.fujifilm.com/about/technology/super_ccd/index.html, 2011.
- [6] www.itl.nist.gov/iad/humanid/feret/feret_master.html, 2011.
- [7] S. Baker and T. Kanade, "Limits on Super-Resolution and How to Break Them," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 9, pp. 1167-1183, Sept. 2002.
- [8] D.F. Barbe, *Charge-Coupled Devices*. Springer-Verlag, 1980.
- [9] M. Ben-Ezra, Z. Lin, and B. Wilburn, "Penrose Pixels: Super-Resolution in the Detector Layout Domain," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 1-8, 2007.
- [10] M. Ben-Ezra, A. Zomet, and S. Nayar, "Video Super-Resolution Using Controlled Subpixel Detector Shifts," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 27, no. 6, pp. 977-987, June 2005.
- [11] S. Borman and R. Stevenson, "Spatial Resolution Enhancement of Low-Resolution Image Sequences: A Comprehensive Review with Directions for Future Research," technical report, Univ. of Notre Dame, 1998.
- [12] T. Chen, P. Catrysse, A.E. Gamal, and B. Wandell, "How Small Should Pixel Size Be?" *Proc. SPIE*, pp. 451-459, 2000.
- [13] D. Daly, *Microlens Arrays*. Taylor and Francis Inc., 2001.
- [14] M. Elad and A. Feuer, "Restoration of Single Super-Resolution Image from Several Blurred, Noisy and Down-Sampled Measured Images," *IEEE Trans. Image Processing*, vol. 6, no. 12, pp. 1646-1658, Dec. 1997.
- [15] S. Farsiu, D. Robinson, M. Elad, and P. Milanfar, "Advances and Challenges in Super-Resolution," *Int'l J. Imaging Systems and Technology*, vol. 14, no. 2, pp. 47-57, 2004.
- [16] R. Fergus, A. Torralba, and W. Freeman, "Random Lens Imaging," *MIT Computer Science and Artificial Intelligence Laboratory TR*, vol. 58, p. 1, 2006.
- [17] W. Freeman and E. Pasztor, "Learning Low-Level Vision," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 1182-1189, 1999.
- [18] G.H. Golub and C.F.V. Loan, *Matrix Computations*, third ed. The John Hopkins Univ. Press, 1996.
- [19] B. Grunbaum and G. Shephard, *Tilings and Patterns*. Freeman, 1987.
- [20] R. Hardie, "A Fast Image Super-Resolution Algorithm Using an Adaptive Wiener Filter," *IEEE Trans. Image Processing*, vol. 16, no. 12, pp. 2953-2964, Dec. 2007.
- [21] R. Hardie, K. Barnard, and E. Amstrong, "Joint Map Registration and High-Resolution Image Estimation Using a Sequence of Undersampled Images," *IEEE Trans. Image Processing*, vol. 6, no. 12, pp. 1621-1633, Dec. 1997.
- [22] R.A. Horn and C.R. Johnson, *Matrix Analysis*, vol. 1. Cambridge Univ. Press, 1985.
- [23] M. Irani and S. Peleg, "Improving Resolution by Image Restoration," *Computer Vision, Graphics, and Image Processing*, vol. 53, pp. 231-239, 1991.
- [24] C. Lin, H. Yang, and C. Chao, "Hexagonal Microlens Array Modeling and Fabrication Using a Thermal Reflow Process," *J. Micromechanics and Microeng.*, vol. 107, pp. 775-781, 2003.
- [25] Z. Lin and H.-Y. Shum, "Fundamental Limits of Reconstruction-Based Superresolution Algorithms under Local Translation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 26, no. 1, pp. 83-97, Jan. 2004.
- [26] L. Middleton and J. Sivaswamy, "Edge Detection in a Hexagonal-Image Processing Framework," *Image and Vision Computing*, vol. 19, no. 14, pp. 1071-1081, 2001.
- [27] S. Park, M. Park, and M. Kang, "Super-Resolution Image Reconstruction: A Technical Overview," *IEEE Signal Processing Magazine*, vol. 20, no. 3, pp. 21-36, May 2003.
- [28] A.J. Patti, M.I. Sezan, and A.M. Tekalp, "Superresolution Video Reconstruction with Arbitrary Sampling Lattices and Nonzero Aperture Time," *IEEE Trans. Image Processing*, vol. 6, no. 8, pp. 1064-1076, Aug. 1997.
- [29] J.B. Pendry, "Negative Refraction Makes a Perfect Lens," *Physical Rev. Letters*, vol. 18, no. 85, pp. 3966-3969, 2000.
- [30] L. Pickup, D. Capel, S. Roberts, and A. Zisserman, "Bayesian Image Super-Resolution, Continued," *Advances in Neural Information Processing Systems*, vol. 19, pp. 1089-1096, 2007.
- [31] I. Shcherback and O. Yadid-Pecht, "Photoresponse Analysis and Pixel Shape Optimization for CMOS Active Pixel Sensors," *IEEE Trans. Electron Devices*, vol. 50, no. 1, pp. 12-18, Jan. 2003.
- [32] T. Sibley and S. Wagon, "Rhombic Penrose Tilings Can Be 3-Colored," *The Am. Math. Monthly*, vol. 107, no. 3, pp. 251-252, 2000.
- [33] D. Takhar, J. Laska, M. Wakin, M. Duarte, D. Baron, S. Sarvotham, K. Kelly, and R. Baraniuk, "A New Compressive Imaging Camera Architecture Using Optical-Domain Compression," *Proc. IS&T/SPIE Computational Imaging IV*, 2006.
- [34] M. Tappen, B. Russell, and W. Freeman, "Exploiting the Sparse Derivative Prior for Super-Resolution and Image Demosaicing," *Proc. Third Int'l Workshop Statistical and Computational Theories of Vision*, 2003.
- [35] W.-Y. Zhao and H.S. Sawhney, "Is Super-Resolution with Optical Flow Feasible," *Proc. European Conf. Computer Vision*, vol. 1, pp. 599-613, 2002.



Moshe Ben-Ezra received the BSc, MSc, and PhD degrees in computer science from the Hebrew University of Jerusalem in 1994, 1996, and 2000, respectively. He was a research scientist at Columbia University from 2002 until 2004 and a member of the technical staff at Siemens corporate research from 2005 until 2007. Since 2007, he has been a lead researcher at Microsoft Research Asia at Beijing. His research interests are in computer vision, with an emphasis on hardware and optics. He is a member of the IEEE.




Zhouchen Lin received the PhD degree in applied mathematics from Peking University in 2000. He is a lead researcher at Visual Computing Group, Microsoft Research Asia. He is now a guest professor at Beijing Jiaotong University, Southeast University, and Shanghai Jiaotong University. He is also a guest researcher at the Institute of Computing Technology, Chinese Academy of Sciences. His research interests include machine learning, computer

vision, numerical computation, image processing, computer graphics, and pattern recognition. He is a senior member of the IEEE.



Bennett Wilburn received the PhD degree in electrical engineering from Stanford University in 2005. He started his engineering career as a VLSI designer working on digital and mixed-signal circuit design for microprocessors at Hewlett Packard. For his thesis work, he designed custom CMOS cameras for a scalable video camera array and devised high-performance imaging methods using the 100-camera system. His recent work has focused on shape

and appearance capture for dynamic scenes. The work in this paper was completed during his five years as a researcher in the Visual Computing group at Microsoft Research Asia in Beijing, China. He is currently working at a computational photography startup company. He is a member of the IEEE.



Wei Zhang received the BEng degree in electronic engineering from Tsinghua University, Beijing, in 2007, and the MPhil degree in information engineering from the Chinese University of Hong Kong in 2009. He is currently working toward the PhD degree in the Department of Information Engineering at the Chinese University of Hong Kong. His research interests include machine learning, computer vision, and image processing.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**