

Penrose Pixels

Super-Resolution in the Detector Layout Domain

Moshe Ben-Ezra, Zhouchen Lin, and Bennett Wilburn

Microsoft Research Asia

{mosheb, zhoulun, bwilburn}@microsoft.com

Abstract

We present a novel approach to reconstruction based super-resolution that explicitly models the detector’s pixel layout. Pixels in our model can vary in shape and size, and there may be gaps between adjacent pixels. Furthermore, their layout can be periodic as well as aperiodic, such as Penrose tiling or a biological retina. We also present 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. Our method can be applied equally well to either periodic or aperiodic pixel tiling.

Through analysis and extensive testing using synthetic and real images, we show that our approach outperforms existing reconstruction based algorithms for regular pixel arrays. We obtain significantly better results using aperiodic pixel layouts. As an interesting example, we apply our method to a retina-like pixel structure modeled by a centroidal Voronoi tessellation. We demonstrate that, in principle, this structure is better for super-resolution than the regular pixel array used in today’s sensors.

1 Introduction

Recent research in super-resolution (SR) have raised significant doubts regarding the usability of reconstruction based super-resolution algorithms (RBA [4]) in the real world. Baker and Kanade [4] showed that the condition number of the linear system and the volume of solutions grow fast with the increment of the magnification factor. Lin and Shum [18] provided a comprehensive analysis of RBA and showed that the effective magnification factor can be at most 5.7. Zhao and Sawhney [24] showed that even achieving proper alignment of local patches for SR is questionable.

To overcome these limitations a different approach to RBA must be taken. As noted by Baker and Kanade [4], RBA can be decoupled into two parts: deblurring of the optical blur and resolution enhancement. Since multiple images taken at *small camera displacements* provide little or no additional information with respect to the optical blur, the first part 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 non-trivial at best.

In this paper we focus on the second aspect of RBA: detector resolution enhancement using multiple images. Optical deblurring can later be applied to the result provided that the

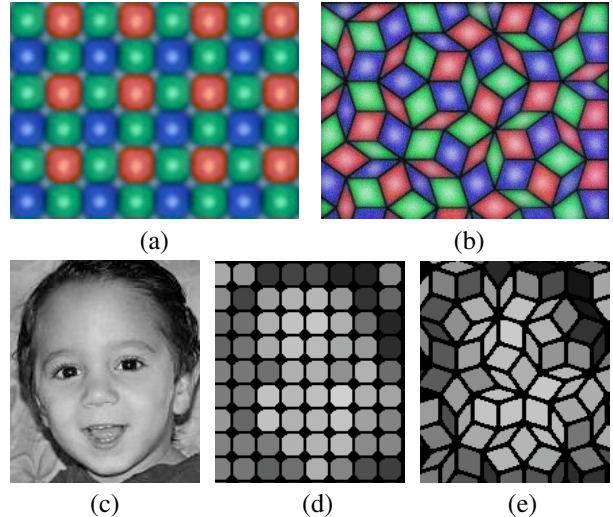


Figure 1. **Regular pixel layout and Penrose pixel layout on the detector plane.** (a) A microscopic view of a Sony 1/3” sensor (part). The shape of the pixels as viewed through the microlens array and color filter array, as well as the gaps between pixels, are clearly visible. (b) A hypothetical aperiodic Penrose pixel layout. (c) An illustration of an image at the sensor surface (irradiance). (d) Spatial integration for the conventional layout. (e) Spatial integration for the Penrose layout.

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 (380–400nm). This is true for high-quality sensors with large pixels (9μ to 15μ) as well as ones with very small pixels (2μ to 4μ). Moreover, sensor technology advances slower than may be expected [22], while physics is already exploring the feasibility of a “perfect lens” using materials with negative indexes of refraction [21]. Therefore, there is a significant need for resolution enhancement at the sensor level.

1.1 Related Work

Roughly speaking, SR algorithms can be categorized into four classes [7, 19, 11]. 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 dealias the LRIs by utilizing the phase difference among the LRIs. Learning-based algorithms (e.g. [12, 4]) 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

$$\mathbf{L} = \mathbf{P}\mathbf{H} + \mathbf{E}, \quad (1)$$

where \mathbf{L} is the column vector of the irradiance of all low-resolution pixels (LRPs) considered, \mathbf{H} is the vector of the irradiance of the HRI, \mathbf{P} gives the weights of the high-resolution pixels (HRPs) in order to obtain the irradiance of the corresponding LRPs, and \mathbf{E} is the noise. To solve (1), various methods, such as maximum *a posteriori* (MAP) [15, 10], regularized maximum likelihood (ML) [10], projection onto convex sets (POCS) [20], and iterative back-projection [16], have been proposed to solve for the HRI.

In all previous work, the LRPs appear on the left hand side of the system (1) and the LRIs are all rectangular regular grids, with square pixels. Based on such a configuration, both the practice and theoretical analysis [4, 18] have shown that the magnification factor is limited to a relatively small number.

1.2 Our Contributions

We find that the magnification factor can be much larger if the aforementioned two conventions of traditional RBAs are changed. First, rather than using the LRPs directly for the left hand side of (1), we instead upsample the LRIs to the high-resolution grid to *match the detector's layout* as shown in Fig. 1(a),(d). 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). Due to the presence of noise, such an underlying treatment results in very different numerical consequence (This is analogous to the preconditioning techniques [13] for solving linear systems.). Second, rather than using a regular grid for the LRIs, we instead use *irregular* pixel layouts for the detector, resulting in LRIs with irregular pixel shapes. The irregular layout helps produce 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, the effective displacement is modulo pixel size, which limits the number of different displacements that are at least ϵ apart).

Moreover, our model's layout, either regular or irregular, does not require that the LRPs fill the whole detector's plane. We specifically model the gaps between physical pixels as *null* values, which better match the information that is *really* acquired by the sensor. In contrast, a sharp detail at the gap between pixels will be wrongly assumed to not exist using the conventional model.

We also propose a novel error back-projection algorithm that iteratively recovers the super-resolved image for arbitrary pixel layouts, either regular or irregular. Our algorithm, though very simple, provides better results than the conventional back-projection algorithms do.




Figure 2. **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 non null HRP values to form the value of its underlying LRP. Resampling is the composition of downsampling and upsampling.

We also discuss other aspects and potential benefits of Penrose tiling based sensors, which are not directly related to super-resolution, and the feasibility of manufacturing Penrose tiling based pixel layouts and matching microlens arrays. We believe that Penrose tiling sensor is feasible with current technology.

2 Penrose Tiling

Penrose tiling is an aperiodic tiling of the plane presented by R. Penrose in 1973 [14]. Fig. 1(b) shows the rhombus Penrose tiling, which consists of two rhombuses, each placed at five different orientations by specific rules [14]. 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 area. Unlike regular tiling, Penrose tiling has no translational symmetry - it never repeats itself exactly¹. 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. In contrast, a regular square tiling forces the maximal delta between different displacements in x and y to be at most $\frac{1}{M}$, where M is the linear magnification factor. The rhombus Penrose tiling shown in Fig. 1(b) is a good candidate for hardware color sensor realization because it is 3-colorable [23] and has simple tiles. This is the primary reason we selected this particular aperiodic tiling.

3 Our Model and Algorithm for SR

We aim at obtaining the best possible results for real photographic systems. Similar to the Jitter-Camera [6], 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, where M is the linear

¹On the infinite plane. But any finite portion of the Penrose tiling can repeat infinitely many times [14].

magnification factor². The shape of LRPs can also be different from each other and gaps between pixels are allowed. As in [5, 4, 18], 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, however as later shown in this section, they can easily be relaxed.

3.1 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 as 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 have an intermediate HRI. It is this intermediate HRI, not the original LRI, that is involved in the computations that follows. As shown in Fig. 2, upsampling is done by placing a regular high resolution pixel grid over the actual shape of the low resolution pixels and then associating HRPs to LRPs. HRPs that are not associated to any LRP are assigned the value *null* to differentiate them from the value zero. The assumption on the uniformity of the pixels can be relaxed at this stage by multiplying the intermediate HRI with a weight mask to compensate for any intra-pixel non uniformities. As we assume that the motion is fronto-parallel translational, and that the displacements between images equal $\frac{1}{M}$, 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 needs to be done 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. 2), 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 an integral over a pixel area, and the upsampling / resampling operator as a flood-fill of the pixel area with this 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 the super-resolved image is the correct one.

3.2 Error Back-Projection Algorithm

Our super-resolution algorithm is a variant of the well-known error back-projection super-resolution algorithm [16]. Unlike the traditional algorithm that downsamples the images into a low resolution array, our algorithm is performed entirely on the high-resolution grid. Using the concepts in the previous

subsection, we summarize our algorithm as follows:

Algorithm 1

Inputs:

- Lo_1, \dots, Lo_{M^2} : Low resolution images ($\sim N^2$ pixels)
- T_1, \dots, T_{M^2} : Transformations for registering the LRIs.
- $M \in \mathcal{N}$: Magnification factor.
- G : Sensor layout map.

Output:

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

Processing:

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

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 up-sample stage. Our upsampling operator $\uparrow_{T, G}$ preserves sharp edges between pixels at the high-resolution grid whereas the conventional algorithm applies the blur kernel globally. Hence our upsampling operator better preserves the high frequencies in the HRI. Also note that if warping is required, it is performed on the intermediate HRI after the upsampling.

4 Analysis

From the linearity of the upsampling and the resampling operators in Algorithm 1, we see that every intermediate HRI I_i aligned to the high-resolution grid is connected to the groundtruth image Su via a matrix A_i :

$$I_i = A_i \cdot Su + n_i, \quad (2)$$

where n_i is the noise from Lo_i . The matrix A_i is exactly the representation of the operator $\downarrow_{T_i, G}$ in Algorithm 1. Therefore, the iteration 3(a) in Algorithm 1 can be written as:

$$Su^{n+1} = Su^n + \frac{1}{M^2} \sum_{i=1}^{M^2} (A_i \cdot Su + n_i - A_i \cdot Su^n), \quad (3)$$

which can be rewritten as:

$$Su^{n+1} - Su^n = (I - \bar{A})(Su^n - Su) + \bar{n}, \quad (4)$$

where $\bar{A} = \frac{1}{M^2} \sum_{i=1}^{M^2} A_i$ and $\bar{n} = \frac{1}{M^2} \sum_{i=1}^{M^2} n_i$. So

$$Su^n - Su = (I - \bar{A})^n (Su^0 - Su) + \left[\sum_{k=0}^{n-1} (I - \bar{A})^k \right] \bar{n}, \quad (5)$$

Assuming that the spectral radius of $I - \bar{A}$ is less than 1, then $\lim_{n \rightarrow \infty} (I - \bar{A})^n = 0$ and \bar{A} is non-singular with $\bar{A}^{-1} =$

$\sum_{k=0}^{\infty} (I - \bar{A})^k$. Then from (5) we have that:

$$\lim_{n \rightarrow \infty} Su^n = Su + \bar{A}^{-1} \bar{n}. \quad (6)$$

² The meaning of the magnification factor for irregular pixel layouts is ambiguous. However, as we focus on quasi-uniform pixel layouts, the magnification factor can still be roughly estimated as the increase in the number of pixels.

4.1 Implications from Error Analysis

From (6), we may expect that the iterations result in a super-resolved image which deviates from the ground truth by $\bar{A}^{-1}\bar{n}$. 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 (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 can choose an appropriate pixel layout so that the norm of \bar{A}^{-1} is small, then the deviation can be effectively controlled regardless of the mean of the noise (note that $\|\lim_{n \rightarrow \infty} Su^n - Su\| = \|\bar{A}^{-1}\bar{n}\| \leq \|\bar{A}^{-1}\| \|\bar{n}\|$).

As $\|\bar{A}^{-1}\|$ is large when \bar{A} is close to be singular, we should choose an appropriate detector's pixel layout such that \bar{A} is far from singular.

According to the above analysis, we must choose pixel layouts such that there are more linearly independent equations in the system (2). 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 [18] also showed that if five LRPs cover the same set of HRPs, then their equation set must be linearly dependent. These indicate that using regular (and square) tiling usually results in insufficient number of independent equation set. To overcome this difficulty, we try to change the regular tiling to other kinds of tilings. An intuition is to use aperiodic tilings.

5 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. As we do not have an actual Penrose pixel sensor, our second experiment strives to be as close to a real world conditions as possible. We first captured a sequence of high-resolution real images (each with its own unique noise characteristics) and then integrated pixel values to simulate a Penrose image. The last experiment is a start-to-finish real image super-resolution test.

5.1 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. 3 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.

We then added Poisson noise (which better models real noise) to the input images. Fig. 4 shows the super-resolution result for the "face" image using additive Poisson noise with mean = 5 and 10 grey levels, followed by 8-bit quantization. Unlike the zero mean quantization error, the non-zero mean Poisson noise significantly degrades the quality of the results. The results can be improved by using many more images than




Figure 4. Noise evaluation test using regular tiling. The top two rows show the result of super-resolution with different magnification factors and Poisson noise of mean 5 (top) and 10 (middle) grey levels. All computations were run for several thousands of iterations or until convergence. The amplification of noise is quite clear and the results are very different from those with zero-mean quantization error. The bottom row shows the results of super-resolution with magnification of x8 and noise mean of 10 running for 1000 iterations, with different numbers of input images (256 to 1600). Results are much better than that obtained by using the minimum 64 input images (the boxed image).

the theoretical minimum requirement, as shown in the bottom row of Fig. 4.

5.2 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. 5 are clearly better than the results obtained with the regular grid, shown in Figs. 3, 4.

To better quantify the results, we used a concentric test target having variable spatial frequency contrast³. 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. 7 shows that our algorithm improves the linear magnification by 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. Figure 8 compares the RMS error as a function of number of images for regular and Penrose tiling respectively. The magnification factor was 8 and the same algorithm (Algorithm-1)

³Real lens' contrast declines as the spatial frequency increases.




Figure 3. **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. The numbers in parentheses are the RMS errors.




Figure 5. **Penrose Pixels super-resolution results.** (top) input images for magnification factor of 8 and 16. (middle) 8-bit quantization result. (bottom) 8-bit quantization, and Poisson noise results. Noise mean equals 5 (using 625 images), and 10 (using 1600 images), respectively.

was applied to both layouts. One can see that while the regular layout improved slightly when over-constrained, the Penrose layout improved by over 4 times. It is interesting to see that the regular layout was actually better when the system was severely under-constrained.

For our last simulation example, we compared our algorithm to an externally obtained result of [15, 4] using an image from




Figure 6. (a) Original image. (b) result of super-resolution using [15] (image taken from [4]). (c) our result using a regular pixel layout. (d) our result using a Penrose Pixels layout.

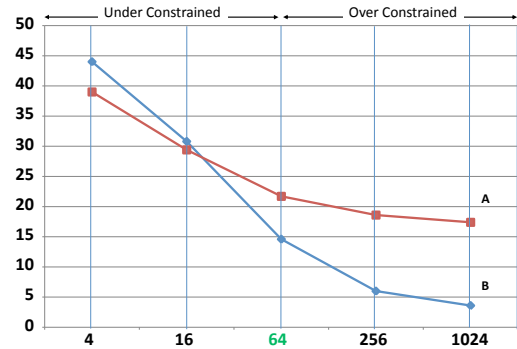


Figure 8. RMS error vs. number of images. (A): regular layout with square pixels. (B): Penrose layout. The Penrose layout clearly better utilizes the additional images.

the FERT database [3]. In Fig. 6 the improvement from our approach is clearly visible.

5.3 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




Figure 7. **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)




Figure 9. **Our algorithm applied to real images (each with its own noise) with simulated pixel integration.** Top: input images for the regular and Penrose pixel layout with magnifications of 8 and 16. Bottom: results of our super-resolution algorithm applied to both regular and Penrose pixel layouts.

each image with the map G , and then quantizing the result. Thus, the resulting LRIs had unique noise due to sensor noise, quantization, and JPEG compression. This process is very similar (though noisier) to pixel binning done at the analog level. As with real sensors’ binning, 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. By capturing more images than the required minimum, we also reduce the effect of slight misalignments. Fig. 9 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.

5.4 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 1(a). Using the lens resolution and pixel size and shape, we created a sensor model for x5 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 image⁴. Fig. 10(a) shows one of the input images and a magnified insert. Fig. 10(b) shows an actual intermediate HRI. The black dots are the null values at the corners of each pixel. Fig. 10(c) 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.

6 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 sim-

⁴This is possible in a controlled environment such as the “Jitter Camera” [6], 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.

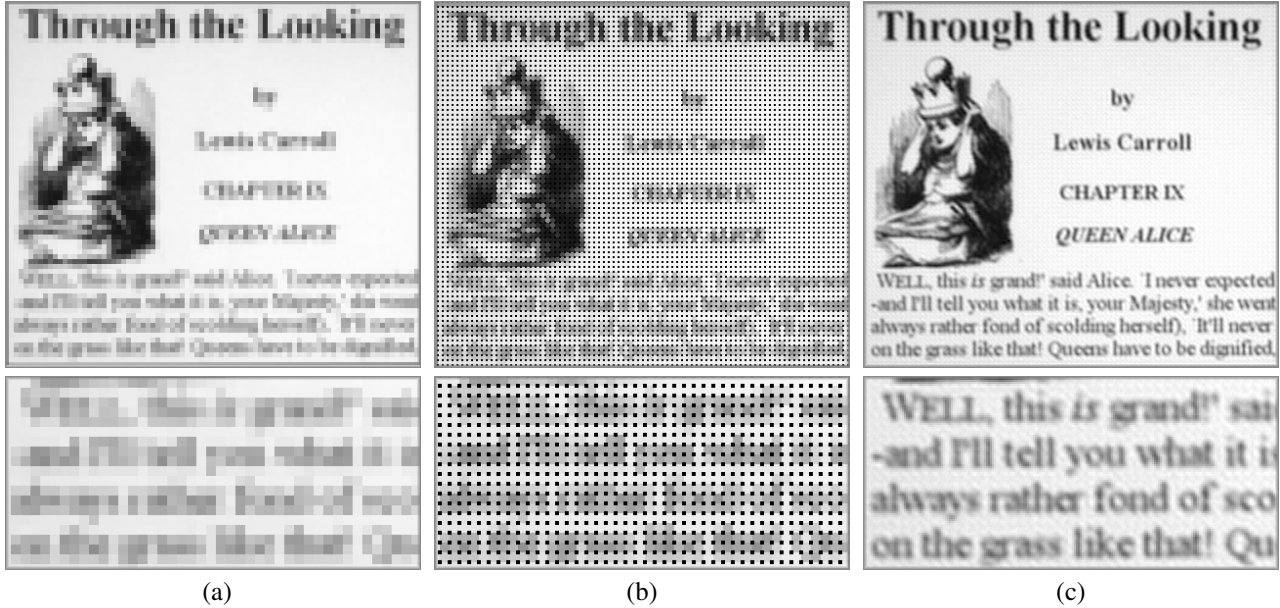


Figure 10. **Real image super-resolution result.** Left: real image captured by the camera (enlarged). Middle: input view for the super-resolution algorithm using our CCD model. The little black dots are gaps between pixels and have null value. Right: super-resolution result. Notice the details at the magnified view, in particular the dots above the letter ‘i’ and in the exclamation marks.

plicity. Further research is needed to determine which tiling, if any, performs best.

Another potentially interesting application of our approach is researching super-resolution in biological vision systems. It is well known that the acuity of many biological vision systems (including the human eye) exceeds the resolution implied by the size of their sensory cells. This phenomena is commonly known as “hyper-resolution”. Though we do not claim that the eye’s hyper-resolution *actually* operates in a similar way to reconstruction based super-resolution, we do believe that our method can estimate the upper limit for the resolving power of biological retinal structures. This can provide an *objective* tool for comparing the potential visual acuity of different animals.

As an example, Fig. 11(a) shows the cell structure of a human fovea. We can see that the structure is nearly hexagonal, but definitely irregular. We modeled this structure by a centroidal Voronoi diagram shown in Fig. 11(b), and we applied our algorithm to the “face” image. The resulting super-resolution image with a magnification factor of roughly 18 is shown in Fig. 11(e). This image looks better and has lower RMS error than the result of the regular pixel layout with a magnification factor of 16. This suggests that the retinal irregular structure theoretically has better resolving power than the regular structure.

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 run on top of the silicon substrate con-




Figure 11. **Super-resolution using a centroidal Voronoi grid.** (a) cross section of a human eye’s fovea. (b) our centroidal Voronoi model with magnification factor of approx. 18 (c) ground truth image. (d) a low-resolution image. (e) reconstructed image RMS error is 9.4.

taining 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 that light enters from the back of the chip. The metal layers are now underneath the photodetectors, so they do not block light. This technology is becoming more appealing as pixel physi-

cal dimensions shrink. Researchers at SONY, for example, recently demonstrated a high-sensitivity, four megapixel color CMOS image sensor with $3.45\mu\text{m}$ square pixels and back-illumination [17].

With no concerns about occluding pixels, Penrose pixel routing becomes much simpler. Generally, in CMOS sensors, each row of pixels shares a signal called a wordline, and each column shares a bitline. When a wordline is asserted, that row of pixels drives their data (a voltage) onto their bitlines to be read out. The wiring challenge is to connect each pixel to a unique wordline/bitline combination. Power supplies and other control signals can run in parallel to these wires. Slightly increasing the density of wordlines and bitlines beyond the theoretical minimum makes this task easy.

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 [2]. We also require microlenses with shapes that match the diamond-shaped pixels. Such microlens arrays can be produced using melting photoresist [9] in a similar way to hexagonal microlens array production [8].

Each of the two pixel shapes occur in five different orientations, so a maximum of only ten unique pixel designs would be necessary. Assuming 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. Thus, we believe it is possible to produce a Penrose pixel image sensor using existing proven technologies.

7 Conclusion

We present a novel approach to super-resolution based on 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, and brings 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 [1] for use with super-resolution. Even a modest 4x linear magnification would turn an 8MP camera into a 128MP one for stationary and possibly moving [6] scenes, without changing the field of view.

Acknowledgments

Penrose tiling postscript program courtesy of Björn Samuelsson.

References

- [1] www.dpreview.com/reviews/minoltadimagea1/. 8
- [2] www.fujifilm.com/about/technology/super_ccd/index.html. 8

- [3] www.itl.nist.gov/iad/humanid/feret/feret_master.html. 5
- [4] S. Baker and T. Kanade. Limits on super-resolution and how to break them. *IEEE PAMI*, 24(9):1167–1183, 2002. 1, 2, 3, 5
- [5] D. F. Barbe. *Charge-Coupled Devices*. Springer-Verlag, 1980. 3
- [6] M. Ben-Ezra, A. Zomet, and S. Nayar. Video super-resolution using controlled subpixel detector shifts. *IEEE PAMI*, 27(6):977–987, 2005. 2, 6, 8
- [7] S. Borman and R. Stevenson. Spatial resolution enhancement of low-resolution image sequences: A comprehensive review with directions for future research. Technical report, University of Notre Dame, 1998. 1
- [8] H. Y. C.P. Lin and C. Chao. Hexagonal microlens array modeling and fabrication using a thermal reflow process. *JOURNAL OF MICROMECHANICS AND MICROENGINEERING*, (107):775C781, 2003. 8
- [9] D. Daly. *Microlens Arrays*. Taylor and Francis Inc., 29W 35 St. New York, NY 10001. 8
- [10] M. Elad and A. Feuer. Restoration of single super-resolution image from several blurred, noisy and down-sampled measured images. *IEEE IP*, 6(12):1646–1658, 1997. 2
- [11] S. Farsiu, D. Robinson, M. Elad, and P. Milanfar. Advances and challenges in super-resolution. *Int'l J. Imaging Systems and Technology*, 14(2):47–57, 2004. 1
- [12] W. Freeman and E. Pasztor. Learning low-level vision. In *Proc. ICCV*, pages 1182–1189, 1999. 1
- [13] G. H. Golub and C. F. V. Loan. *Matrix Computations (3rd edition)*. The John Hopkins University Press, Baltimore, Maryland, 1996. 2
- [14] B. Grunbaum and G. Shephard. *Tilings and Patterns*. Freeman, 1987. 2
- [15] R. Hardie, K. Barnard, and E. Amstrong. Joint map registration and high-resolution image estimation using a sequence of undersampled images. *IEEE IP*, 6(12):1621–1633, 1997. 2, 5
- [16] M. Irani and S. Peleg. Improving resolution by image restoration. *Computer Vision, Graphics, and Image Processing*, 53:231–239, 1991. 2, 3
- [17] S. Iwabuchi, Y. Maruyama, Y. Ohgishi, M. Muramatsu, N. Karasawa, and T. Hirayama. A back-illuminated high-sensitivity small-pixel color cmos image sensor with flexible layout of metal wiring. In *International Solid State Circuits Conference (ISSCC)*, San Francisco, CA. 8
- [18] Z. Lin and H.-Y. Shum. Fundamental limits of reconstruction-based superresolution algorithms under local translation. *IEEE PAMI*, 26(1):83–97, 2004. 1, 2, 3, 4
- [19] S. Park, M. Park, and M. Kang. Super-resolution image reconstruction: A technical overview. *IEEE Signal Processing Magazine*, 20:21–36, 2003. 1
- [20] A. J. Patti, M. I. Sezan, and A. M. Tekalp. Superresolution video reconstruction with arbitrary sampling lattices and nonzero aperture time. *IEEE IP*, 6(8):1064–1076, 1997. 2
- [21] J. B. Pendry. Negative refraction makes a perfect lens. *Phys. Rev. Lett.*, 18(85):3966 – 3969, 2000. 1
- [22] A. E. G. T. Chen, P. Catrysse, and B. Wandell. How small should pixel size be? *Proc. SPIE*, (3965):451C459, 2000. 1
- [23] S. W. T. Sibley. Rhombic penrose tilings can be 3-colored. *The American Mathematical Monthly*, 107(3):251–252, 2000. 2
- [24] W.-Y. Zhao and H. S. Sawhney. Is super-resolution with optical flow feasible? 1