

High Quality Alias Free Image Rotation

Charles B. Owen

Fillia Makedon

Dartmouth Experimental Visualization Laboratory

Dartmouth College

6211 Sudikoff Labs, Hanover, NH 03755

cowen@cs.dartmouth.edu

Abstract

This paper presents new algorithms for the rotation of images. The primary design criteria for these algorithms is very high quality. Common methods for image rotation, including convolutional and separable approaches, are examined and shown to exhibit significant high frequency aliasing problems. A new resampling filter design methodology is presented which minimizes the problem for conventional convolution-based image rotation. The paper also presents a new separable image rotation algorithm which exhibits improved performance in term of reduction in artifacts and an efficient $O(N^2 \log N)$ running time.

1. Introduction

Rotation is a very basic image processing operation. Many applications, such as radiology and photographic analysis, require very high quality image rotation. The goal of the work presented here is efficient image rotation of the highest possible quality.

An image is defined, for the purpose of this paper, as a discrete sampling of a continuous two-dimensional function onto a square lattice. If an image is to be rotated by θ degrees, the sample points of the rotated image fall between the original image lattice points. The sampling theorem provides a basis for the reconstruction of the continuous function which the image samples represent (This assumes that the original image was band-limited by the Nyquist frequency in both dimensions. If this is not the case the aliased components have become a part of the continuous image the samples now represent, so this issue can be disregarded in this discussion) [2]. The sampling theorem also provides a basis for the resampling of this continuous function on the new, rotated, lattice. The combination of

reconstruction and resampling allows for high accuracy rotation of images.

In realizable systems the reconstruction and resampling filters can be combined. This leads to computations for the new sample points based on discrete convolution. The computed sample is a weighted sum of the samples in a finite neighborhood of the lattice point, as described in [5, 11] in the more general context of image warping. The work presented herein will give a basis for the design of correct convolution filter kernels for image rotation.

Convolution-based rotation is very costly computationally. Quality is dependent upon the size of the interpolation kernel [8, 9], increasing as the kernel size increases with maximum quality when the kernel size is as large as the image. For an $N \times N$ image and an $M \times M$ convolution kernel M^2 multiply operations must be performed for each of the N^2 pixels of the image. For the highest quality, $M \rightarrow N$, tending towards a running time of $O(N^4)$. This problem was recognized in early image rotation papers such as [10, 3]. These and later papers show that if the rotation can be decomposed into operations only on rows and columns of the image the operations can employ 1-D convolution rather than 2-D convolution. These decompositions are referred to as *separable image rotations*.

Two-pass separable image rotation reduces rotation to two scale and skew operations. The first operation reduces image resolution, thereby decreasing image resolution in intermediate images [3, 10, 5]. This decrease in resolution is referred to as *bottle-necking* and results in loss of high frequency content. Three-pass decompositions, as presented in [4, 9], implement rotation using only skew operations and no scaling. While significantly higher quality than the two-pass techniques, these approaches still introduce significant aliasing, as demonstrated in this paper.




Figure 1. Rotated spectrum. The spectrum of an image has been rotated 45 degrees. The corners of the spectrum (shown dimmed here) extend beyond the Nyquist frequency, folding back into the sides as shown.

2. Convolution-Based Rotation

Convolution-based rotation is the simplest form of rotation. No intermediate images are required and the rotation algorithm is trivial (given a filter kernel). The process of resampling is described in detail in [7]. Briefly, the continuous image is reconstructed from the sample points and resampled at the rotated lattice points. The resampling process must low-pass filter the rotated image to remove frequencies beyond the Nyquist frequency in each dimension. Since no scaling occurs, at first glance it would appear that frequency characteristics are unchanged since the image is not magnified or minified as discussed in [11]. However, rotation of an image induces an equivalent rotation of the spectrum of the image[6]. As shown in figure 1, spectral content in the corners of an image can rotate beyond the Nyquist frequency.

Assuming frequency characteristics such that no aliasing occurs, only an interpolation filter would be required. Interpolation kernels approximate a box filter as closely as possible. Indeed, [9, 11] spend considerable time discussing alternative filter kernels. However, this approach causes corner aliasing due to the fact that the frequency response of the reconstructed continuous image will be a rotated square.

A solution to this problem is to account for this fold-back when resampling the image. A filter kernel is constructed which not only interpolates, but also filters off the components which would be rotated beyond the Nyquist frequency bounds. A simple method for constructing this filter kernel is to compute the inverse Fourier transform of the multiplication of two box filters $A(u, v)$ and $B(u, v)$ such that $A(u, v)$ represents the interpolation filter and $B(u, v)$ is a box filter rotated by $-\theta$. This is simply computed using an inverse

FFT.

An ideal filter of this form would require infinite support. Since infinite support is not practical, the filter must be truncated, which induces Gibbs oscillations[1]. In our experimentation we have employed a simple Hamming window with good results.

While simple to understand, convolution-based image rotation has problems, the worst being its computational complexity. Given an $N \times N$ image and an $M \times M$ filter kernel, rotation requires $O(N^2 M^2)$ time. Another serious problem is that a filter kernel is computed for a given resampling location relative to adjacent pixels. For any given offset $\delta_{x,y}$ a new filter kernel must be computed. The filter kernel computation is costly so most practical implementations compute a subset of filter kernels and store these in a lookup table[11]. This approach still requires computing a considerable number of filters (which may consume a large amount of space) and the filters represent quantized approximations of the true sampling points.

3. Separable Rotation

Separable image rotation decomposes the rotation into operations which only involve rows or columns of the image. A three-step decomposition consisting of only skew operations as described in [5, 9]:

$$R(\theta) = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} = \begin{bmatrix} 1 & -\tan\theta/2 \\ 0 & 1 \end{bmatrix} \times \begin{bmatrix} 1 & 0 \\ \sin\theta & 1 \end{bmatrix} \times \begin{bmatrix} 1 & -\tan\theta/2 \\ 0 & 1 \end{bmatrix} \quad (1)$$

The operations necessary to perform the skew consist only of row or column shifts. In this case, a fixed shift δ_r is applied to each row (or column). Figure 2 illustrates the rotation of an image using this three-pass approach.

The three pass approach induces interesting spectral effects which contribute to errors in the resulting image. In order to understand what occurs when an image undergoes three part separable rotation, the following theorem is required.

Theorem 1 *If $f(x, y)$ has the Fourier transform $F(u, v)$, then $f(x - \sigma y, y)$ has the Fourier transform $F(u, v + \sigma u)$.*

Proof: The proof of this theorem is straight-forward using the substitutions $\phi(s, t) = s + \sigma t$ for x and $\psi(s, t) = t$ for y :

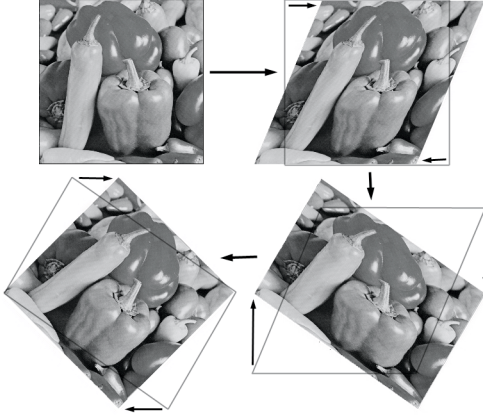


Figure 2. Three pass separable image rotation. The top left image is the original image. Steps proceed in a clockwise order through the figure. Arrows illustrate the skew operation.

$$\begin{aligned}
 \mathcal{F}f(x - \sigma y, y) &= \\
 &= \int \int_{-\infty}^{\infty} f(x - \sigma y, y) e^{-j2\pi(ux+vy)} dx dy \\
 &= \int \int_{-\infty}^{\infty} f(s, t) e^{-j2\pi(us+(v+u\sigma)t)} \begin{vmatrix} 1 & \sigma \\ 0 & 1 \end{vmatrix} ds dt \\
 &= F(u, v + \sigma u) \tag{2}
 \end{aligned}$$

Corollary 1 If $f(x, y)$ has the Fourier transform $F(u, v)$, then the inverse Fourier transform of $F(u, v + \sigma u)$ is $f(x - \sigma y, y)$.

This theorem and corollary states that a skew of σ in one dimension of an image is equivalent to skew of the *spectrum* of that image by $-\sigma$ in the other dimension. Hence, each step in a separable image rotation is also inducing skew in the spectrum of the image. The problem is that the spectrum of a square image is also square (without loss of generality, images will be assumed to be square in this discussion). The spectrum following the skew will no longer be square, but rather a parallelogram which will not fit in the square bounds enforced by the sampling theorem.

An ideal skew operation of a sampled image is equivalent to sampling the continuous image after skew. The various approaches to skew realization work this way in that they compute the interpolated point between adjacent pixels necessary to perform the skew. The image is assumed band limited prior to the skew, so the points of the parallelogram after the skew induce




Figure 3. Aliasing after skew.




Figure 4. Spectral shuffling. Only half of the spectrum areas are labeled. The remaining areas are symmetrical.

aliasing in the image as illustrated in figure 3. This aliasing occurs in every step of the rotation process. Figure 4 illustrates the spectral shuffling which occurs after a 3-pass rotation of 45 degrees.

It is expected that some components of the spectrum will need to be masked off if the image is to be alias-free. This is due to the fact that a diamond spectrum is now to be fit into a square shaped hole. This could be easily dealt with using post-filtering. However, the two grayed slivers in figure 4 are within the valid spectral bounds and are an obvious error condition, so simple post-filtering of the image will not correct these aliased components. Basically, this is an incorrect shuffling of the spectrum.

Figure 5 illustrates an in-place rotation of an image.



Figure 5. Spatial shuffling.

The sliver in question is quite obvious. Spatially, the solution seems obvious: simply pad the image width to avoid the loss of content in the first rotation step. The width must be padded by a multiplicative factor of $1 + \tan^2(\theta/2)$ to avoid loss. However, padding an image spatially expands the image spectrum. Hence, an approach might be to take an FFT of the padded image, pad it again by a factor of $1 + \tan^2(\theta/2)$. Given that FFT operations will be used, the first pad may double the size of the image to reach a multiple of two. If the second padding must be inverted using an inverse FFT, again a double in size will occur. The final result will be 16 times as much image data.

The algorithm presented below provides a solution to this problem. The image padding is performed twice as mentioned before, but only one dimension is padded at each step.

4. Rotation Algorithm

Let N be the image size and θ a rotation angle such that $-90^\circ \leq \theta \leq 90^\circ$. N' is a usable FFT size greater than $N + N \tan^2(\theta/2)$. N'' is a usable FFT size greater than $N \times 2$.

1. The image is widened to N' wide and $N''/2$ tall, padding around the image with black.

2. An FFT is performed on all rows of the image.

3. The rows are multiplied by $\exp(-j2\pi n\delta/N')$, where δ represents the appropriate shift of the row necessary to implement a skew of $-\tan\theta/2$.

If an inverse FFT were performed on the rows at this point, the result would be an image after the first skew operation. The algorithm needs to now perform a two-dimensional FFT on the image. Since the Fourier transform is separable and one part of the transform is already done, there is no need to inverse transform the rows.

4. An FFT is performed on the columns. At this point the data is equivalent to a 2D FFT of the entire widened image including the first skew step.

5. The image is doubled to N'' vertically. The first widening preserved spatial information which would have been lost. This second widening preserves spectral information. The widening process replicates image content above and below rather than padding with zeroed spectrum. The aliased components remain, but the spectrum widening copies back what was shifted into aliasing before.

6. An inverse FFT is performed on the rows.

7. The *rows* are multiplied by $\exp(j2\pi n\delta/N')$, where δ represents the appropriate shift of the *column* necessary to implement a skew of $-2\sin\theta$. The doubling is required because the image height is halved

(doubling spectrum height halves spatial height).

8. An FFT is performed on the image rows.

This operation is equivalent to the second skew in the three pass rotation. That skew is a column skew in the spatial domain. The algorithm performs a negative row skew in the spectral domain, which is equivalent.

9. Window functions are applied to the spectral representation to remove the components which will be aliased. (These can be removed since they do not mix with existing frequencies, but rather shuffle to holes left by the skew process). The spectrum to be retained at this point is in the shape of a parallelogram with the corners clipped off.

10. An inverse FFT is performed on the columns.

11. The third pass skew is implemented by multiplying the *columns* by $\exp(j2\pi n\delta/N'')$, where δ represents the appropriate shift of the *row* necessary to implement a skew of $(\tan\theta/2)/2$.

At this point the intermediate image is equivalent to a dual padded final image after FFTs have been applied to rows only. The vertical content is band limited due to the earlier masking process.

11. An inverse FFT is performed on the image rows.

12. The image is decimated vertically by discarding every other sample, then cropped to its original size.

This algorithm is fully separable and has a running time of $O(N^2 \log N)$ for an $N \times N$ image. It also uses the exact same number of forward and inverse 1-D FFTs (three sets) as the SEP3-sinc algorithm described as the highest quality in [9]. This complexity is similar to three 2-D FFT operations (1.5 times the time to forward and inverse FFT the entire image).

This algorithm is based on the use of the Fourier *shift property* described in [2] and the shift technique described in [9]. To shift a row by δ , the FFT of the row is taken, the transform coefficients are multiplied by $\exp(-j2\pi n\delta/N)$, effectively adjusting the phase of the components. Then the inverse FFT is taken. In many cases the algorithm leaves the result in the FFT state. The algorithm also takes advantage of Theorem 1 to decrease the number of FFT operations.

The Fourier shift property is well known and discussed in most basic texts. It is trivial to show that, given a simple sign change in the proof of the shift property, an inverse shift property can be shown. The inverse property is used to shift a line of the spectrum. A shift of δ of the spectrum can be computed by simply performing an inverse FFT, multiplying each sample s_n by $\exp(j2\pi n\delta/N)$, and reapplying the FFT.

Given a real image, the algorithm produces a real image as a result. This is due to the fact that, though complex operations are performed on the image and much of the rotation is done in the complex domain, all

intermediate images are equivalent to those produced by the spatial only three pass approach up to the removal of the aliased components.

5. Experimental Results

Testing rotation algorithms is not a simple proposition. Ground truth data for non-trivial rotation of an image usually does not exist. One approach that has been taken is to rotate by some angle, then rotate back. This is actually a poor test because the scrambling of the spectrum in the rotation is *exactly* undone by the reverse rotation. The test approach utilized in this paper is to perform two successive 45 degree rotations and compare to an image rotated 90 degrees using a trivial pixel remapping. (Even this technique is not ideal, since the intermediate 45 degree image rotation will clip corners of the spectrum. However, this problem only serves to inflate errors on the tests. Hence, actual quality is somewhat higher.) A more ideal test technique would utilize synthetic images which can be generated at any rotation angle. Such a test is under development, but was not ready for presentation with this paper.

Three images were used for testing: Peppers, Lena, and Photographer. For each image an ideal rotation of 90 degrees was generated and compared to rotations using the three pass separable rotation of [9] (referred to as SEP3-sinc in that paper) and the three pass separable rotation algorithm in this paper. The images were padded with black for SEP3-sinc. The technique detailed in this paper includes its own padding technique. Results for these images are listed in Table 1. For each image the peak signal-to-noise ratio and the mean squared error for the image is illustrated.

These results exhibit significant decreases in distortion relative to the previous approach. The remaining distortion can be attributed to several factors. The error results were computed on 8 bit binary images rather than the intermediate floating point image in the algorithm, so there is some expected quantization error yielding a ceiling of 54dB PSNR. Also, some numerical inaccuracies due to roundoff in computations are to be expected. Finally, the implementation used for this test used a simple rectangular filter for masking alias components. This filter will exhibit Gibbs oscillations. A windowing filter such as a Hamming filter will probably yield even better results.

6. Summary

This paper presents solutions very high quality image rotation including the definition of alias free convo-

Test Image	SEP3-sinc		New	
	PSNR	MSE	PSNR	MSE
Lena	40.58	5.69	44.52	2.29
Peppers	40.14	6.30	45.49	1.84
Photographer	40.21	6.20	44.41	2.36

Table 1. Experimental Results

lution rotation kernels as well as a new three-pass separable image rotation algorithm. The paper also demonstrates aliasing issues of common approaches to image rotation including both convolution and conventional three pass separable image rotation, demonstrating significant aliasing of high frequencies in those techniques. The more powerful FFT based separable algorithm presented herein is alias free and achieves a running time of $O(N^2 \log N)$ for an $N \times N$ image.

References

- [1] A. Antoniou. *Digital Filters: Analysis and Design*. McGraw-Hill, New York, 1979.
- [2] R. N. Bracewell. *The Fourier Transform and Its Applications*. McGraw-Hill, New York, second edition, 1986.
- [3] E. Catmull and A. R. Smith. 3-d transformations of images in scanline order. In J. J. Thomas, editor, *SIGGRAPH'80 Conference Proceedings*, pages 279–285. Association for Computing Machinery, 1980.
- [4] P.-E. Danielsson and M. Hamnerin. High-accuracy rotation of images. *CVGIP: Graphical Models and Image Processing*, 54(4):340–344, 1992.
- [5] D. Fraser. Comparison at high spatial frequencies of two-pass and one-pass geometric transformation algorithms. *Computer Vision, Graphics, and Image Processing*, 46:267–283, 1989.
- [6] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1992.
- [7] P. S. Heckbert. Fundamentals of texture mapping and image warping. Masters thesis, University of California, Berkeley, 1989.
- [8] M. Unser, M. A. Neimark, and C. Lee. Affine transformations of images: A least squares approach. In *IEEE Int. Conference on Image Processing*, volume III, pages 558–561, Austin, TX, 1994.
- [9] M. Unser, P. Thévenaz, and L. Yaroslavsky. Convolution-based interpolation for fast high-quality rotation of images. *IEEE Transactions on Image Processing*, 4(10):1371–1381, 1995.
- [10] C. F. R. Weiman. Continuous anti-aliased rotation and zoom of raster images. In J. J. Thomas, editor, *SIGGRAPH'80 Conference Proceedings*, pages 286–293. Association for Computing Machinery, 1980.
- [11] G. Wolberg. *Digital Image Warping*. IEEE Computer Society Press, 1990.