Quasi-Circular Rotation Invariance in Image Denoising

Yi Wan and Robert D. Nowak*
Department of Electrical and Computer Engineering
Rice University, USA

Abstract

This paper studies a new method for wavelet-based image denoising which is translation invariant (TI) and *rotation invariant (RI). These invariances are crucial in image denoising and, more generally, may play important roles in image modeling. In contrast to other approximately RI methods, like the steerable pyramid, our new method employs standard separable wavelet bases in conjunction with a pseudo-circular image rotation. This scheme does not involve interpolation and hence the observation model (likelihood function) is invariant under this rotation. The superiority of our new method with respect to existing TI (non-RI) techniques is supported by experiments.*

1 Introduction

A major difficulty in wavelet-based image denoising (and wavelet-based image modeling in general) is the lack of translation invariance (TI) and rotation invariance (RI) associated with separable wavelet bases. TI can be obtained fairly easily using the frame formed by all possible translations of the wavelet basis, but RI is more difficult to achieve simply because images are usually sampled on a rectangular lattice. It is well-known that TI denoising methods reduce blocky artifacts and in general produce better image denoising results than non-TI methods [1, 2]. However, TI methods still tend to favor certain features that align well with the wavelet basis functions. In the case of Haar wavelets, for example, these features are approximately vertical or horizontal edges. For a slant edge with slope of, say, 30 degrees, or a circular disk, it is easy to see that their shape and orientation will remain unchanged for any translation operation. Hence, TI methods are not well-matched to the above mentioned situations which are very common in natural images. The geometry of natural images can be dealt with more effectively by designing more complex, often redundant, image representations. The shiftable multiscale transform developed in [3, 4] represents one effort in this direction. The shiftable multiscale transform (here shiftable refers to both translation and rotation) essentially involves a standard image rotation. Unfortunately, standard rotation is not compatible with rectilinear image sampling schemes, and involves pixel interpolation. Here we propose a new rotation scheme that is better suited to square image lattices and offers certain advantages.

In this paper we develop a new translation and rotation invariant (TRI) denoising method which uses a quasi-circular image rotation scheme. It makes use of the l_∞ distance norm, rather than the Euclidean norm underlying standard image rotation, and hence we call this an l_∞-rotation. The main advantages of the TRI method include the following. First, for sampled images the l_∞-rotation is equivalent (in the sense of norms) to regular image rotation. Secondly, this method preserves the simple statistical structure of the image observation model (*e.g.,* Gaussian white noise or Poisson counts). Thirdly, our method does a better job at noise removal and artifact suppression in wavelet-based shrinkage/threshold denoising than regular and TI methods. And lastly, the TRI method doesn't involve any interpolation or extra memory in the implementation. In fact, computation-wise, each rotation costs about the same as a translation. All these lead to a very clean approach and efficient implementation scheme.

The rest of paper is organized as the following. We first briefly review the TI method in Section 2, then develop the RI method and explore the properties of the l_∞-rotation and its relationship to regular rotation in Section 3. Combining TI and RI methods leads to a new class of TRI methods, which are studied in Section 4. Throughout the development, emphasis is on the application to image denoising. Experimental results demonstrate the potential of TRI wavelet-based methods. Possible applications to other image analysis and modeling applications are also discussed. Conclusions are made in Section 5.

*This work was supported by the National Science Foundation, grant no. MIP–9701692

2 TI Wavelet-Based Denoising

The TI method was proposed to reduce blocky artifacts in wavelet thresholding image denoising schemes [1, 2]. It uses all possible shifts of an (orthogonal) wavelet basis to form a frame. The frame is hence redundant and the method is also called redundant DWT (RDWT). For more background on RDWT, see, for example, Burrus *et al* [5]. The standard TI approach consists of the following three steps. Translate the noisy data to get a new image; transform the translated image into wavelet domain, apply a certain wavelet shrinkage/thresholding method such as hard or soft thresholding, and inverse-transform back to the space domain; then inverse-translate the denoised image back to the original position. Finally, the average of the multiple (one for each translation) denoised images is taken as the final result.

Suppose that we observe and image s with additive Gaussian white noise:

$$y = s + n \,, \tag{1}$$

where $n \sim \mathcal{N}(0, \sigma^2 I)$, with $\sigma^2 I$ denoting the diagonal noise covariance. Then, the TI denoising procedure is summarized as

$$\widehat{s} = \frac{1}{N} \sum_n T_{-n} W^{-1} D_\delta \left(W T_n y \right) \,, \tag{2}$$

where N is the total number of possible translations, W represents the DWT, T_n is the translation operations for vector $n = (n_x, n_y)$, and D_δ is a shrinkage/thresholding operator with threshold δ. Very efficient schemes exist for implementing the TI method with complexity of $O(N \log N)$. In all experiments shown in this paper we employ the Haar wavelet basis and the hard thresholding rule [1] with threshold level set to $\sqrt{3}\sigma$, where σ is the standard deviation of the additive noise.

3 RI Denoising via l_∞-Rotation

Standard image rotation involves the interpolation of pixels and often brings with it added computational overhead and intractability in modeling. We develop here a non-rigid image rotation called l_∞-rotation. This operation can be viewed in the following three equivalent ways. Assume that the center of the image is the origin of the spatial coordinate system. Let (x, y) denote the Cartesian coordinates of any pixel point in the image.

1. This view is similar to that of regular rotation. Here each point is defined in polar coordinates (r, θ), but the distance is measured with the l_∞ norm, i.e.,
$$r = \max(|x|, |y|)$$
instead of the usual Euclidean norm $\sqrt{x^2 + y^2}$ (hence the name l_∞-rotation). Then a rotation of angle α is defined as
$$T_r(\alpha) \;:\; (r, \theta) \longrightarrow (r, \theta + \alpha) \tag{3}$$

2. In this view, we use the usual Euclidean norm and $T_r(\alpha)$ is defined as
$$T_r(\alpha) \;:\; (r, \theta) \longrightarrow (r', \theta + \alpha) \tag{4}$$
where
$$r' = r \, \frac{\max\{|\cos\theta|, \; |\sin\theta|\}}{\max\{|\cos(\theta + \alpha)|, \; |\sin(\theta + \alpha)|\}} \tag{5}$$
Eq. 5 is an equivalent way of saying that the l_∞ norm remains the same after the l_∞-rotation.

3. This view is an intuitive one. Here an image is considered as consisting of square tracks of pixels. For example, around the center there is the smallest square track of only 4 pixels; the outermost track consists of $4M - 4$ pixels where M is the number of pixels in each row or column (assume square images). In an l_∞-rotation, each pixel is only allowed to move along its own square track.

An example of an l_∞-rotation is shown in Figure 1. Unlike regular rigid transforms such as translation or usual rotation, l_∞-rotation also involves some "twisting" and "bending".

(a) (b)

Figure 1: *(a) Image. (b) 45 degree l_∞-rotation.*

Now we can develop an RI denoising scheme analogous to the TI algorithm in Eq. 2. The RI denoising estimate is

$$\widehat{s} = \frac{1}{N} \sum_n R_{-\theta_n} W^{-1} D_\delta \left(W R_{\theta_n} y \right) \tag{6}$$

where N is the total number of possible rotations and R_{θ_n} represents the l_∞-rotation of angle θ_n. For an $M \times M$ image, there are $N = 4M - 4$ unique l_∞-rotations.

A major advantage of the l_∞-rotation, compared to standard rotations, is that it preserves (probabilistic) image observation models. For example, given the observation model in Eq. 1, we have

$$y|s \ \sim \ \mathcal{N}(s, \sigma^2 I)\,. \tag{7}$$

Applying a rotation,

$$R_{\theta_n} y|s \ \sim \ \mathcal{N}(R_{\theta_n} s, \sigma^2 I)\,; \tag{8}$$

the covariance is still diagonal and we have uncorrelated errors. Standard image rotations would produce a non-diagonal covariance structure. In the case of Poisson observation models for photon-limited imaging, it is easy to show that the Poisson nature of the data is also preserved under l_∞-rotations; standard rotations would completely destroy the Poisson model.

Images that undergo an l_∞-rotation will usually appear to varying degrees distorted for rotations that are not multiples of 90 degrees. But this achieves the purpose of aligning the non-vertical or non-horizontal edges with a separable wavelet basis, and is at least comparable to standard rotations. In Figure 2, we depict several common shapes undergoing l_∞-rotations. Some general observations can be made regarding the properties of l_∞-rotation with comparison with regular rotation. First, we see that a straight line segment is mapped to piece-wise straight line under l_∞-rotation, while it remains a straight line segment under usual rotation. Secondly, a circle is distorted into a different shape under l_∞-rotation. Even though it looks less regular, as far as commonly used wavelets (such as Haar) are concerned, this shape may be even prefered because it is closer to the shape of a square. Also note that just as a circle centered at the origin remains the same under any regular rotation, a square centered at the origin also remains the same under any l_∞-rotation. We also see that a square is rotated into a polygon under l_∞-rotation. Again in this case even though usual rotation keeps a square shape, it does not have apparent advantage in fitting the rotated square with wavelet bases. Lastly and most importantly, a line passing through the origin remains a line after any l_∞ rotation.

The implementation of RI method is just as simple as that of the TI method because interpolation is not required. In Figure 3, the RI method is compared with the TI method (both using Haar wavelets) and shows better denoising effect. This confirms the

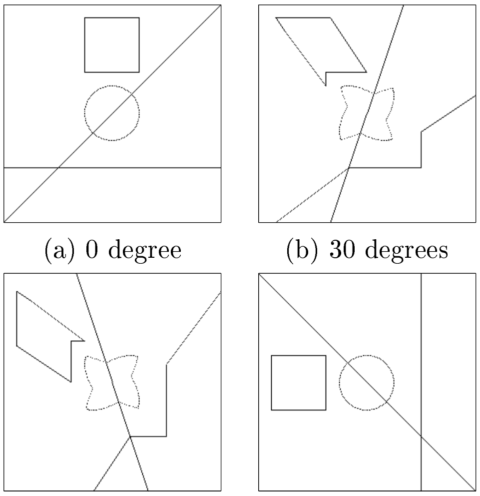

(a) 0 degree (b) 30 degrees (c) 60 degrees (d) 90 degrees

Figure 2: l_∞-*rotation of some common shapes.*

intuition that rotation does a better job in aligning image components with wavelets than translation since singularities are usually not vertically or horizontally aligned.

So we now have two approaches that provide desirable invariances in wavelet-based estimation. However, the TI method lacks necessary flexibility to deal with commonly occured singularities along curves and slant edges. The RI model overcomes this, but it is not translation invariant and tends to distort image structures execpt cases such as lines passing through the origin. In the next section, we present a hybrid method that is both TI and RI.

4 TRI Wavelet-Based Denoising

By combining Eq. 2 and Eq. 6 in the previous two sections, we have a translation and rotation invariant denoising scheme given by

$$\widehat{s} = \frac{1}{N_t} \sum_n T_{-n} \frac{1}{N_r} \sum_m R_{-\theta_m} W^{-1} D_\delta \left(W R_{\theta_m} T_n y \right) \tag{9}$$

where N_t and N_r are the total numbers of translations and rotations considered, respectively.

Note that just like the case with regular rotation, since T_n and R_θ are not commutative, RTI method above is different from the TRI method (*i.e.,* first rotate, then translate). Moreover, we could employ repeated applications of the translation and rotation operator (*e.g.,* TRTI) to generate other possibilities. For notational simplicity we call all cases the TRI method, but in this paper we focus on the form given in Eq. 9. We next prove the following proposition which makes the TRI method especially useful.

Proposition 1 *Under any l_∞-rotation, a line segment passing through the origin is mapped to another line segment with length scaled by a factor between $1/\sqrt{2}$ and $\sqrt{2}$.*

Proof. We use the second view point on the l_∞-rotation discussed in Section 3. Here the l_∞-rotation of angle α can be viewed as a product of two transformations. i.e.,

$$R_\alpha = D \cdot R_{l_2}(\alpha)$$

where R_α is the l_∞-rotation of angle α, $R_{l_2}(\alpha)$ is the regular (Euclidean) rotation operator of angle α and D is a scaling factor.

Now given any line segment l passing through the origin, $R_{l_2}(\alpha)\, l$ is just another line segment of equal length through the origin with the angle increment of α. The scaling factor D is given by Eq. 5. Recall that for any value of θ,

$$\max\{|\cos\theta|, |\sin\theta|\} \in \left[\frac{1}{\sqrt{2}}, 1\right]$$

It follows that the scaling factor $\frac{r'}{r} \in \left[\frac{1}{\sqrt{2}}, \sqrt{2}\right]$
□

The above proposition shows that if we use the TRI method, then an edge anywhere in an image can be first translated to a position so as to pass through the origin, then rotated to any orientation. In the TRTI approach, it is even possible to rotate any edge to any orientation at the original position locally. Furthermore, because the l_∞-rotation also induces different scaling factors in different parts of the image, even curved edges can be "bent" into different shapes as shown in Figure 2 having a more rectangular structure, which are thus more likely to align better with wavelets. In Figure 4, the TRI method (with Haar wavelets) in Eq. 9 is used which considers all possible translations and 15 rotations (uniformly spaced over 0 to 90 degrees). This example shows that the TRI method out-performs both TI both RI methods.

5 Conclusions and Extensions

In this paper, a new non-rigid image rotation scheme is employed to develop the RI and TRI denoising methods for better noise reduction over the existing algorithms. The latter outperforms translation invariant image denoising by doing a better job in aligning singularities with the underlying wavelet basis functions. Unlike other methods such as the shiftable filters, l_∞-rotation faithfully retains the desirable properties of separable wavelet bases, such as regularity, across all scales. The simple structure of these methods not only allows a very efficient implementation, but may also offer an attractive framework for more complicated image models. For example, the l_∞-rotation can be directly extended to voxels in 3-D data. Moreover, in general, the l_∞-rotations can also be applied directly to the Fourier and cosine packet domains.

We also note that the l_∞-rotation described here is an extension to the "pseudo-circular" rotations studied by Geman and Koloydenko [6]. In that work, the authors identify image equivalence classes based on this type of rotation as a way to support "geometric" image invariances. They argue that image statistics should be fairly invariant to such rotations. Their work attests to the importance of l_∞-rotation in image modeling and denoising.

References

[1] R. Coifman and D. Donoho, "Translation invariant denoising," in *Lecture Notes in Statistics: Wavelets and Statistics*, New York: Springer-Verlag, pp. 125-150, 1995.

[2] M. Lang and H. Guo and J. E. Odegard and C. S. Burrus and R. O. Wells, "Noise reduction using an undecimated discrete wavelet transform," *IEEE Signal Processing Letters*, vol. 3, no. 1, pp. 10-12, 1996.

[3] W. T. Freeman and E. H. Adelson "Steerable filters". *Topical Meeting on Image Understanding and Machine Vision*, vol. 14. Optical Society of America, June 1989. Technical Digest Series.

[4] E. P. Simoncelli, W. T. Freeman, Edward H. Adelson and D. J Heeger, "Shiftable Multi-scale Transforms", *IEEE Trans. Information theory*, vol. 38(2), pp. 587-607, March 1992

[5] C. S. Burrus, R. A. Gopinath and H. Guo, *Introduction to Wavelets and Wavelet Transforms, A Primer*, Prentice Hall, 1998.

[6] D. Geman and A. Koloydenko, "Invariant Statistics and Coding Natural Microimages," Technical Report, Univ. of Mass., 1998.

(a) MSE=1

(b) MSE=0.4892

(c) MSE=0.3148

Figure 3: *TI and RI denoising using Haar wavelets. (a) Noisy image. (b) TI denoising with 16 translations. (c) RI denoising with 15 rotations.*

(a) MSE=0.3148

(b) MSE=0.2918

(c) MSE=0.2747

Figure 4: *RI, TI and RTI denoising using Haar wavelets. (a) RI denoised image with 15 rotations. (b) TI denoised image with all possible translations. (c) RTI denoised image combining (a) and (b)*