PLATELETS FOR MULTISCALE ANALYSIS IN PHOTON-LIMITED IMAGING

Rebecca M. Willett and Robert D. Nowak

Department of Electrical and Computer Engineering
Rice University, MS-366, 6100 Main, Houston, TX, 77005 USA

ABSTRACT

This paper proposes a new multiscale image decomposition based on platelets. Platelets are localized functions at various scales, locations, and orientations that produce piecewise linear image approximations. For smoothness measured in certain Hölder classes, the error of m-term platelet approximations can decay significantly faster than that of m-term approximations in terms of sinusoids, wavelets, or wedgelets. Platelet representations are especially well-suited for the analysis of Poisson data, unlike most other multiscale image representations, and they can be rapidly computed. We propose a platelet-based maximum penalized likelihood criterion that encompasses denoising, deblurring, and tomographic reconstruction.

1. PHOTON-LIMITED IMAGING

In a variety of applications, data are acquired by the detection of (light or higher energy) photons, and often the random nature of photon emission and detection is the dominant source of noise in imaging systems. Such cases are referred to as *photon-limited* imaging applications, since the relatively small number of detected photons is the factor limiting the signal-to-noise ratio. These applications include Positron Emission Tomography (PET), Single Photon Emission Computed Tomography (SPECT), Confocal Microscopy, astronomical imaging, and Infrared (IR) imaging [1].

Wavelet-based methods are powerful tools for image denoising and have been used in applications of this nature; however, most wavelet-based approaches are based on Gaussian approximations to the Poisson likelihood, and these approximations can be especially inaccurate in the context of photon-limited imaging applications because of the low count levels. The Haar wavelet system is the only exception; it does provide a tractable multiscale analysis framework for Poisson data. Haar analyses are restrictive, however, in that they can only result in piecewise constant approximations and cannot approximate smooth boundaries or curves and edges well.

R. Nowak was partially supported by the National Science Foundation, grant no. MIP–9701692, the Army Research Office, grant no. DAAD19-99-1-0349, the Office of Naval Research, grant no. N00014-00-1-0390. R. Willett was partially supported by the National Science Foundation Graduate Student Fellowship.

In answer to these difficulties, we propose a new multiscale image representation based on atoms called platelets. Platelets are localized atoms at various locations, scales and orientations that can produce highly accurate, piecewise linear approximations to images consisting of smooth regions separated by smooth boundaries. At the heart of fast, platelet-based, maximum penalized likelihood methods for image estimation and reconstruction lies a recursive partitioning scheme based on multiscale likelihood factorizations which, unlike conventional wavelet decompositions, are very well suited to applications with Poisson data. Platelet-based techniques for image denoising and estimation significantly outperform even clairvoyant wavelet thresholding techniques. Because platelet decompositions of Poisson distributed images are tractable and computationally efficient, expectation-maximization type algorithms enhanced with platelet techniques can easily be used to reconstruct images in an ill-conditioned inverse problem setting.

2. PLATELET APPROXIMATIONS

One of the key ideas underlying multiresolution analysis is that of recursive partitions, which yield sequences of nested image partitions. Recursive partitions (RP) in general are interesting because they allow for important extensions of classical Haar multiscale analysis. The wedgelet partition [2] is based on a dyadic, square recursive partition which allows for non-square, "wedge-shaped" partitions only at the final level of the partition. That is, a wedgelet partition is based on a recursive dyadic square partition of the image in which the final partition "cells" are allowed to terminate with a wedge instead of a square.

In general, an RP can contain smaller / larger dyadic squares or wedges in different regions of the image domain. This spatial variation in the "resolution" or "scale" of the partition enables partitions to "zoom-in" in regions of fine detail and structure and "zoom-out" in smooth regions of the image. Spatially varying partitions can be constructed by beginning with a very fine (pixel-level) partition of the image merging partition cells (squares or wedges) in a bottom-up fashion, according to a data-based metric of fit versus partition complexity.

Instead of approximating the image on each cell of the

partition by a constant, as is done in a Haar or wedgelet analysis, we can approximate it with a planar surface. We define a platelet $f_S(x,y)$ to be a function of the form

$$f_S(x,y) = (A_S x + B_S y + C_S)\ I_S(x,y), \qquad (1)$$

where $A_S, B_S, C_S \in \mathbb{R}$, S is a dyadic square or wedge associated with a terminal node of an RP, and I_S denotes the indicator function on S. Each platelet requires three coefficients, compared with the one coefficient for piecewise constant approximation.

In many applications it is beneficial to have the added flexibility of platelet approximations. Image gradients, or smooth transitions between regions of varied intensities, encode information about light emission or reflection as well as surface geometry. Platelet approximations are ideally suited to photon-limited applications because of the inherent smoothness of the images. In medical imaging applications, radioactive pharmaceuticals diffuse smoothly in regions of homogeneous tissue, resulting in smoothly varying intensities within images of organs, and in astronomical applications, images of celestial bodies are similarly smooth. The goal of these applications, however, is often to detect boundaries of or abnormalities within objects under observation. Such regions would be indicated by sharp changes, *i.e.* edges, in the intensity image.

In the context of photon-limited imaging, therefore, it is natural to assume an image model consisting of smooth regions separated by a smooth boundary. The power of platelets is realized in connection with m-term approximations. Although each platelet has two more parameters per term, for images of sufficient smoothness many fewer platelets than constant blocks are needed to approximate an image to within a certain error [3]. For this analysis, consider images which are Hölder smooth apart from a Hölder smooth boundary over $[0,1]^2$:

Theorem 1 *Consider the class of images*

$$f(x,y) = f_1(x,y)\cdot I_{\{y \geq H(x)\}} + f_2(x,y)\cdot\left(1 - I_{\{y \geq H(x)\}}\right)$$

$\forall (x,y) \in [0,1]^2$ *where* $f_i \in$ *Hölder*$^{\beta,2}(C_\beta)$, $i = 1,2$, *and* $H \in$ *Hölder*$^{\alpha,1}(C_\alpha)$ *with* $\alpha, \beta > 1$. *Suppose that* $2 \leq m \leq 2^J$, *with* $J > 1$. *The squared* L_2 *error of* m*-term,* J*-scale, resolution* δ *platelet approximation to images in this class is less than or equal to* $K_{\alpha,\beta} m^{-\min(\alpha,\beta)} + \delta$, *where* $K_{\alpha,\beta}$ *depends on* C_α *and* C_β.

Theorem 1 (proved in our platelets technical report [3]) shows that for images consisting of smooth regions ($\beta > 1$) separated by smooth boundaries ($\alpha > 1$) m-term platelet approximations may significantly outperform Fourier, wavelet, or wedgelet approximations, which have rates of $O(m^{-\frac{1}{2}})$, $O(m^{-1})$, and $O(m^{-1})$, respectively, for this class. Figure 1 displays such an image (a linear top and quadratic bottom separated by a cubic boundary) and an approximation of it with platelets. Wavelets and Fourier approximations do not perform well on this class of images due to the boundary. Conversely, wedgelets can handle boundaries of this type, but produce piecewise constant approximations and perform poorly in the smoother (but non-constant) regions of images.

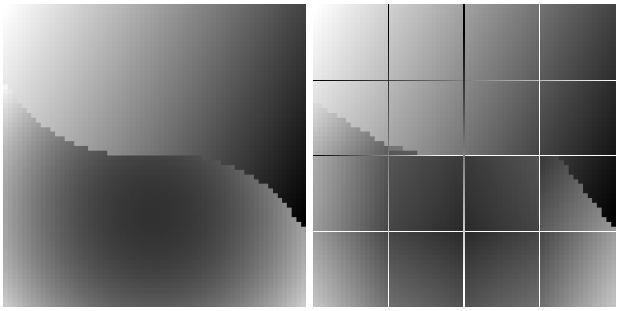

(a) (b)

Fig. 1. Approximation Example. (a) Image of smooth regions separated by a smooth boundary. (b) Rough approximation of (a) with platelets (lines drawn to illustrate platelet boundaries).

3. PLATELET ANALYSIS OF POISSON DATA

Suppose that $x(u,v)$ is a realization of a Poisson process. Underlying this process is an continuous intensity function $\mu(u,v)$, $(u,v) \in [0,1]^2$. Assume that either by choice or perhaps the limitations of measuring instruments, $x(u,v)$ is observed only discretely on the squares (pixels) $S_{m,n}$, $m,n = 0,\ldots,N-1$. It is assumed that the effect of the discretization is to yield an array of count measurements $\boldsymbol{x} \equiv \{x_{m,n}\}_{m,n=0}^{N-1}$, associated with an array of intensity parameters $\boldsymbol{\mu} \equiv \{\mu_{m,n}\}_{m,n=0}^{N-1}$. Each $x_{m,n}$ is simply the number of events in the square $S_{m,n}$ and $\mu_{m,n} \equiv \int_{S_{m,n}} \mu(u,v)$. The likelihood of $\boldsymbol{x}$, given the intensities or probabilities $\boldsymbol{\mu}$, is denoted by $p(\boldsymbol{x}|\boldsymbol{\mu})$.

In earlier work, we introduced a class of *multiscale likelihood factorizations* for Haar multiscale analysis that provided an alternative probabilistic representation (*i.e.*, in addition to that of the original likelihood $p(\boldsymbol{x}|\boldsymbol{\mu})$) of the information in $\boldsymbol{x}$, in a manner indexed by the various location/scale combinations offered by a given recursive partition [4]. The likelihood factorization allows the likelihood of the entire image to be represented in a tree structure in which both likelihoods and parameter penalties of children are inherited by parents. Thus the likelihood factorization serves as a probabilistic analogue of an orthonormal wavelet decomposition of a function. In our technical report [3] we extended the class of multiscale likelihood factorizations to encompass likelihoods generated by platelet analyses. The parameters of the conditional likelihoods play the same role as wavelet coefficients in a conventional wavelet-based multiscale analysis. Using this, it is possible to optimally prune

a recursive partition of the data using a fast algorithm reminiscent of dynamic programming and the CART algorithm [3]. The pruning process is akin to a "keep or kill" wavelet thresholding rule.

The optimal pruning solves a maximum penalized likelihood estimation, wherein the penalization is based on the complexity of the underlying partition. The complexity of a given partition is proportional to the total number of terminal cells. Our goal here is to maximize the penalized likelihood function

$$L_\gamma(\boldsymbol{\mu}) \equiv \log p(\boldsymbol{x} \mid \boldsymbol{\mu}) - \gamma\{\#\boldsymbol{\theta}\}, \tag{2}$$

where $p(\boldsymbol{x} \mid \boldsymbol{\mu})$ denotes a likelihood factorization and $\{\#\boldsymbol{\theta}\}$ is the number of parameters in the vector $\boldsymbol{\theta}$ (one for each constant terminal cell, three for each platelet terminal cell). The constant $\gamma > 0$ is a weight that balances between fidelity to the data (likelihood) and complexity regularization (penalty), which effectively controls the bias-variance tradeoff.

The solution of

$$\begin{aligned} (\widehat{\mathcal{P}}, \widehat{\boldsymbol{\theta}}) &\equiv \arg\max_{\mathcal{P}, \boldsymbol{\theta}} L_\gamma(\boldsymbol{\mu}(\mathcal{P}, \boldsymbol{\theta})) \\ \widehat{\boldsymbol{\mu}} &\equiv \boldsymbol{\mu}(\widehat{\mathcal{P}}, \widehat{\boldsymbol{\theta}}) \end{aligned} \tag{3}$$

is called a maximum penalized likelihood estimator (MPLE). Larger values of γ produce smoother, less complex estimators; smaller values of γ produce more complicated estimators. The best overall performance (as measured by MSE) depends on the choice of γ.

The algorithm is detailed in our technical report [3], and finds the globally optimal solution to (3), as desired. The computational complexity of platelet analysis is not significantly greater than that of wavelet or wedgelet analyses. We bound the computational complexity of performing either an "approximate" or exact MPLE, where the "approximate" estimate uses a (suboptimal) least-squares platelet fit and the exact estimate is obtained by numerically optimizing the (concave) log likelihood function for the most likely platelet fit. The theorem below is proved in [3].

Theorem 2 *(MPLE): A Haar MPLE can be computed in $O(N^2)$ operations, where N^2 is the number of pixels in the image. A wedgelet or "approximate" platelet MPLE can be computed in $O(N^3)$ operations. An exact platelet MPLE can be computed in $O(N^4)$ operations.*

4. PHOTON-LIMITED IMAGE DENOISING AND RECONSTRUCTION

Platelet-based estimation techniques are here compared with wavelets, applied to nuclear medicine data for denoising, and applied to confocal microscopy data for deblurring.

SNR	Platelets	D4 Wavelets	D6 Wavelets
1	0.038	0.041	0.048
10	0.029	0.101	0.130
100	0.039	0.159	0.198

Table 1. $\frac{MSE}{SNR}$ for Denoising with Platelets and Wavelets

4.1. Comparison with Wavelets

In order to compare the performance of platelets and wavelets, we generated Poisson data using Figure 1(a) as a scaled version of the true intensity at resolution 128×128 pixels. For SNRs (average intensities per pixel) of one, ten, and one hundred, seventy-five realizations were generated and denoised with platelets and D4 and D6 wavelets. The wavelet hard threshold level was chosen clairvoyantly for each noisy image to yield the minimum mean square error, and the MPLE penalty $\gamma = \frac{1}{2}\log(\#\text{counts})$ was used for the platelet estimator. The MSEs (averaged over all trials, and normalized by the total intensity) for each SNR are displayed in Table 1. Clearly platelets have a significant advantage over conventional wavelet denoising methods.

4.2. Nuclear Medicine Denoising

Figure 2 (a) depicts an image of the spine obtained from a nuclear medicine study. Functional changes in the bone can be detected using nuclear medicine image before they will show up in X-ray images. Here we applied a technique called "averaging over shifts" or "cycle-spinning". This entails circularly shifting the raw data by a few pixels, denoising, and then shifting the estimate back to its original position. Five shifts in each direction (horizontal and vertical) yielded a total of twenty-five estimates, which were then averaged. This technique often improves denoising and reconstruction results because it reduces the dependence of the estimator on the dyadic partition. The result of the application of our denoising algorithm is displayed in Figure 2 (c). Comparisons with the result of hereditary Haar wavelet denoising (in Figure 2 (b)) demonstrate the superiority of platelet representations of image structures. The improvements in edge clarity are most visible in the spinal column and hips.

Platelet-based techniques can also be used in tomographic reconstruction applications such as PET, as detailed in our technical report [3]. The algorithm is similar to the one used for confocal microscopy image reconstruction in the following section.

4.3. Confocal Microscopy

Confocal microscopy is used to obtain volume images of small fluorescent objects with high spatial resolution. Due to the geometry of these microscopes, a "blurring" is in-

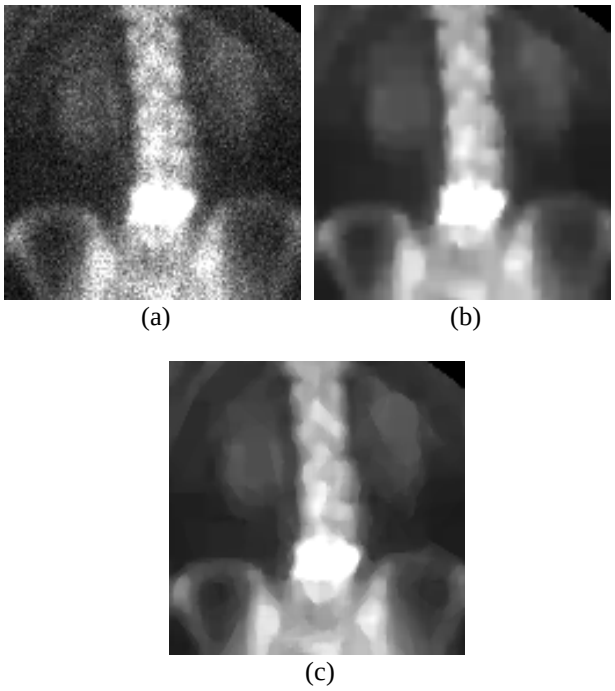

Fig. 2. Denoising in nuclear medicine. (a) "Raw" nuclear medicine spine image (256×256 pixels). (b) Haar-based MPLE, $\gamma = \frac{1}{5}\log(\#\text{counts})$, averaged over 25 shifts. (c) Platelet-based MPLE, $\gamma = \frac{1}{5}\log(\#\text{counts})$, averaged over 25 shifts.

troduced into the measurement process. This distortion of the image is commonly modeled by the convolution of the true image with the point spread function of the microscope. Since the arrival of fluorescence light at the photo multiplier tube can be modeled as a Poisson process, the "de-blurring" and estimation process may be viewed as a Poisson inverse problem well suited to the application of iterative platelet estimation using the EM algorithm, as detailed in [3].

A common use for CFMs is the imaging the dendritic spines of neurons; with this application in mind, we developed the 128×128 phantom image in Figure 3(a). In practice, some regions of the object will be closer to the detector or exhibit more fluorescence than others, which we model in our phantom with image gradients, making this an excellent candidate for platelet analysis. The averaged-over-shifts platelet-based MPLE appears in Figure 3(c). For comparison purposes, we also present the results obtained by using the commonly applied stopped EM-MLE algorithm in Figure 3(b), here stopped at the iteration yielding the minimum MSE. Finally, Figure 3(d) plots the L_2 error of both of the estimates (and that of a Haar-based MLPE) at each iteration. After several iterations the EM-ML estimate worsens considerably with each subsequent iteration. In contrast, both MPLEs converge eliminating the need to choose which iteration is the best stopping point, as done with the EM-MLE. Furthermore, the converged MPLEs exhibit significantly less L_2 error than the best MLE.

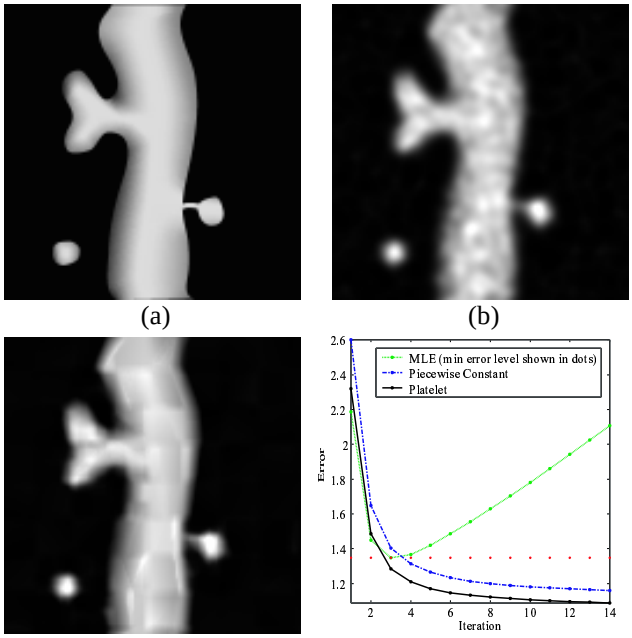

Fig. 3. Confocal microscopy simulation. (a) Phantom (128×128 pixels). (b) Best EM-MLE restoration. (c) Platelet-based MPLE (averaged over 3×3 shifts). In this case, $\gamma = \frac{1}{5}\log(\#\text{counts})$ and convergence was declared when $||\boldsymbol{\mu}^{(i+1)} - \boldsymbol{\mu}^{(i)}||_2 / ||\boldsymbol{\mu}^{(i)}||_2 < 10^{-5}$ (14 iterations in this case). (d) MSE decay by iteration.

5. SUMMARY

This paper introduced the platelet representation for the analysis, denoising, and reconstruction of photon-limited images. Platelets outperform conventional wavelet representations because of their ability to approximate smooth boundaries more efficiently than wavelets. Experimental results with real data from astronomy, confocal microscopy and nuclear medicine demonstrate the effectiveness of platelet-based methods. Astronomical imaging results will appear in the final paper.

6. REFERENCES

[1] D. Synder and M. Miller, *Random Point Processes in Time and Space*, New York: Springer-Verlag, 1991.

[2] D. Donoho, "Wedgelets: Nearly minimax estimation of edges," *Ann. Statist.*, vol. 27, pp. 859 – 897, 1999.

[3] R. Willett and R. Nowak, "Platelets: A multiscale intensity estimation of piecewise linear Poisson processes," Tech. Rep. TREE0105, Rice University, 2001.

[4] E. Kolaczyk and R. Nowak, "A multiresolution analysis for likelihoods: Theory and methods," submitted to Annals of Stat. Available at http://www.ece.rice.edu/ nowak/publications.html.