

IMAGE RESTORATION USING A SPARSE QUADTREE DECOMPOSITION REPRESENTATION

Adam Scholefield and Pier Luigi Dragotti

Communications and Signal Processing Group, Imperial College London.

Email: adam.scholefield@imperial.ac.uk, p.dragotti@imperial.ac.uk

ABSTRACT

Techniques based on sparse and redundant representations are at the heart of many state of the art denoising and deconvolution algorithms. A very sparse representation of piecewise polynomial images can be obtained by using a quadtree decomposition to adaptively select a basis. We have recently exploited this to restore images of this form, however the same model can also provide very good sparse approximations of real world images. In this paper we take advantage of this to develop both image denoising and deconvolution algorithms suitable for real world images. We present results on the cameraman image showing comparable performance with soft thresholding using the undecimated wavelet transform in the denoising case and iterative soft thresholding using the undecimated wavelet transform in the deconvolution case.

Index Terms— Image restoration, piecewise polynomial approximation, quadtrees, sparse matrices.

1. INTRODUCTION

Image restoration is a classic well studied problem that has many practical applications. It is common to assume the following degradation model,

$$y = Hx + e, \quad (1)$$

where y is the noisy blurred image, H is the matrix representing the convolution, x is the desired image and e is additive Gaussian white noise.

As sparsity is not normally present in the image domain directly but in the transform domain, let us first define a linear transformation $x = D\theta$. D is the matrix reconstructing the image from the transform coefficients θ . The columns of D are the basis functions of the approximation space and D can thus be thought of as a dictionary of basis functions. It is expected that the coefficients vector θ will be sparse. To impose this we will usually assume that the probability density function (pdf) of θ is $p(\theta) = K \exp[-\lambda \|\theta\|_p^p]$ (The K is just a constant to ensure that the function is in fact a valid pdf). Using this assumption the maximum a posteriori (MAP)

estimator of x is given by:

$$\begin{aligned} \hat{x} &= D\hat{\theta}, \quad \text{where} \\ \hat{\theta} &= \arg \max_{\theta} p(\theta|H, D, y) \\ &= \arg \max_{\theta} \ln \left[\frac{1}{p(y)} p(y|H, D, \theta) p(\theta) \right] \\ &= \arg \min_{\theta} \|y - HD\theta\|_2^2 + \lambda \|\theta\|_p^p. \end{aligned} \quad (2)$$

The $p(y|H, D, \theta)$ is known from the assumption of the structure of the noise e . The first term of (2) is the data misfit term and the second term imposes sparsity. The choice of λ controls the importance of the two terms and p defines how to enforce sparsity. When $p = 0$ the norm is defined as the number of non zero coefficients of θ : mathematically this is not a norm however it does truly impose the sparsest solution.

The solution of (2) has received much recent research interest. In the trivial case when $H = I$ and D is a unitary transform the cost function is exactly minimised by simple shrinkage (i.e. hard-thresholding for $p = 0$ and soft-thresholding for $p = 1$). When $1 \leq p \leq 2$ the more general case is exactly solved by iteratively applying shrinkage functions. The convergence of this iteration was first proved by Daubechies et al [1] and the technique has produced many image restoration algorithms [2, 3, 4]. Many of these algorithms were proposed well before the proof of global convergence. For a good overview of iterated shrinkage algorithms see [5]. Blumensath and Davies [6] have proved that the iteration converges to a local minimum in the non convex case ($p < 1$).

The choice of transformation D is of course critical to the performance of the restoration method. Wavelets are most commonly used and are the current standard in image processing applications, despite this the quest for sparser representations of images is still receiving much research interest. The main problem with two-dimensional wavelets is that they can only efficiently represent point singularities and not higher order singularities such as edges which are a key part of real world images. Motivated by this, Shukla et al [7] developed a compression algorithm tailor made for piecewise polynomial images. Their algorithm was based around quadtree decomposition and was able to outperform the JPEG2000 standard, due to the sparser representation

achieved by their transformation. Shukla et al also developed a denoising algorithm [8] using the same piecewise polynomial model and rate-distortion cost function as [7]. Recently we proposed image denoising and deconvolution algorithms suitable for piecewise polynomial images [9] using the same model as [7, 8], but using a different cost function not based on the rate which is more appropriate to compression. In this paper we extend these ideas to the restoration of real world images. The rest of the paper is organised as follows: in Section 2.1 the proposed quadtree decomposition algorithm is introduced, then Section 2.2 uses this for image denoising. Section 2.3 develops an iterative deconvolution algorithm using the same model and a surrogate function. We also develop a technique to perturbate from local minima to an approximation with a smaller cost. Section 3 shows the results of both denoising and deconvolution of the camera-man image and deconvolution of a piecewise linear image: the proposed algorithms are compared against wavelet based methods. Finally Section 4 gives conclusions.

2. PROPOSED IMAGE RESTORATION ALGORITHMS

2.1. Piecewise polynomial model using a quadtree decomposition

A quadtree decomposition dyadically partitions an image into a tree structure. The whole image is represented by the root of the tree; the image is then split into four equal quadrants that correspond to the four children of the root node. The splitting process is recursively iterated on each leaf in the tree resulting in many small regions corresponding to the leaves at the deepest depth in the tree. The flexibility in the approach comes from the fact the algorithm chooses which nodes in the tree to use depending on the signal that is being approximated. In the piecewise polynomial model of [7, 8, 9] each image region is approximated by either a polynomial or two polynomials separated by a continuous boundary. Figure 1a shows an example of how this model can approximate a piecewise linear image where all the boundaries are straight edges. For simplicity in the rest of this paper we will assume this simpler model for the boundaries, however the ideas easily extend to the more general cases.

One problem with just using the regions from a dyadic partition is that neighbouring regions can only be jointly represented if they have the same parent. To overcome this we first find the best dyadic regions and then look to join neighbouring regions. This results in the partitions of Figure 1b.

2.2. Denoising algorithm using the piecewise polynomial model

To denoise we assume that the original image is well approximated by a representation that is sparse in the transform domain. To impose this we use the cost function given in (3);




Fig. 1: Example of prune and prune-join partitions

this minimises a tradeoff between the data misfit term and solutions with minimum description length.

$$\hat{\theta} = \arg \min_{\theta} \|y - D(\theta)\|_2^2 + \lambda P(\theta), \quad (3)$$

where $D(\theta)$ is the image representation with coefficients θ (as the decomposition is nonlinear we use $D(\theta)$ rather than $D\theta$). $P(\theta)$ is the function to penalise the description length of the approximation and is defined in the following way:

$$P_i(\theta) = \begin{cases} 2d + 1 & \text{if global piece} \\ 2d_1 + 2d_2 + 2 + \ln(a) & \text{if two pieces} \end{cases},$$

where d, d_1, d_2 are the degrees of the polynomials over the global region or either side of the edge and a is the area of tile i . P_i is the penalty associated with the representation of tile i and the global penalty is simply given by the sum over all tiles, $P(\theta) = \sum_i P_i(\theta)$. The $2d + 1$ term is basically equivalent to a zero norm term, for example a globally linear term would require three basis functions, a constant and a linear function in both the x and y directions. When there is an edge these terms are present for both of the regions. Since the number of possible discrete edges for a particular tile size is proportional to a we use the $\ln(a)$ term to represent this description cost.

2.3. Deconvolution algorithm using the piecewise polynomial model

To use the piecewise polynomial model for deconvolution we simply need to insert the the blurring matrix H into the denoising cost function (3):

$$\hat{\theta} = \arg \min_{\theta} \|y - HD(\theta)\|_2^2 + \lambda P(\theta). \quad (4)$$

Unfortunately the H causes all the basis functions in our transformation to overlap which means that we cannot locally look for the best tile. To solve this we take inspiration from the linear transform case and use a surrogate function and the MM philosophy to decouple these equations, for a good introduction to MM algorithms see [10]. Equations (5) and (6)

show the original cost function C and surrogate cost function C_{sur} respectively.

$$C(D(\theta)) = \|y - HD(\theta)\|_2^2 + \lambda P(\theta) \quad (5)$$

$$C_{sur}(D(\theta) | a) = C(D(\theta)) - \|HD(\theta) - Ha\|_2^2 + \alpha \|D(\theta) - a\|_2^2 \quad (6)$$

It can easily be shown that the surrogate function is a maximiser of $C(D(\theta))$ if $\alpha \geq \|H\|$. I.e.

$$C_{sur}(D(\theta) | a) \geq C(D(\theta)) \quad \forall \theta \quad (7)$$

$$C_{sur}(a | a) = C(a). \quad (8)$$

The surrogate function has the advantage that the $\|HD(\theta)\|_2^2$ terms cancel essentially decoupling the equations:

$$\begin{aligned} C_{sur}(D(\theta) | a) &= \|y\|_2^2 - 2D(\theta)^T H^T y + \|HD(\theta)\|_2^2 \\ &\quad - \|HD(\theta)\|_2^2 + 2D(\theta)^T H^T Ha \\ &\quad - \|Ha\|_2^2 + \alpha \|D(\theta)\|_2^2 + \alpha \|a\|_2^2 \\ &\quad - 2\alpha D(\theta)^T a + \lambda P(\theta) \end{aligned} \quad (9)$$

$$\begin{aligned} \frac{C_{sur}(D(\theta) | a)}{\alpha} &= \left\| a + \frac{H^T}{\alpha} (y - Ha) - D(\theta) \right\|_2^2 + \bar{\lambda} P(\theta) \\ &\quad + \text{terms independent of } \theta. \end{aligned} \quad (10)$$

We can see that the minimisation of the surrogate function is equivalent to minimising the denoising cost function (3) with y replaced with $a + \frac{H^T}{\alpha} (y - Ha)$. The MM approach suggests to thus solve the problem with the iteration:

$$\theta^{i+1} = \text{Denoise} \left(D(\theta^i) + \frac{H^T}{\alpha} (y - HD(\theta^i)) \right). \quad (11)$$

From the inequalities of a maximiser we know that

$$C(D(\theta^{i+1})) \leq C_{sur}(D(\theta^{i+1}) | D(\theta^i)) \quad (12)$$

$$C_{sur}(D(\theta^i) | D(\theta^i)) = C(D(\theta^i)). \quad (13)$$

Therefore if the denoising algorithm guarantees that $C_{sur}(D(\theta^{i+1}) | D(\theta^i)) \leq C_{sur}(D(\theta^i) | D(\theta^i))$ then the sequence will always be decreasing. As the previously introduced denoising algorithm only approximately solves (3) then this is not guaranteed in its current state. To guarantee this we use an update algorithm which starts looking for the best approximation from the current representation θ^i .

When this iteration has converged we look to escape from the local minimum by updating a single tile or pruning four children to their parent. This is possible by noticing that although all the basis functions in (5) overlap preventing a closed form solution, it is still possible to update only one whilst fixing all the others. We try this for all the individual tiles and update the single representation that results in the greatest decrease in cost.

As with the denoising case, we first assume the simpler prune only model and introduce deconvolution joining algorithms when the pruning algorithms have converged.



(a) Original



(b) Degraded, PSNR=16.08 dB



(c) Reconstructed with wavelets, PSNR=26.62 dB



(d) Reconstructed with proposed algorithm, PSNR=26.99 dB

Fig. 2: Denoising of cameraman image (512×512)

3. EXPERIMENTAL RESULTS

In the following experiments the proposed algorithms were implemented with polynomials of maximum degree 1.

Figure 2 shows the result of denoising the cameraman image that has been degraded by additive white Gaussian noise of standard deviation 40 resulting in a PSNR of 16.08 dB. The proposed denoising algorithm is compared against wavelet soft thresholding using the undecimated wavelet transform. The Daubechies 4 tap filter was used to a depth of 6. In this case the performance of the techniques are comparable however our algorithm suffers in areas with high texture such as the grass. This suggests that adding a different tile model to our representation may significantly improve our algorithm in areas that are not well approximated by piecewise polynomials.

Figures 3 and 4 show the results of deconvolving a piecewise polynomial image and a real world image. The piecewise polynomial image was degraded by blurring with a 7 by 7 quadratic spline followed by additive white Gaussian noise of standard deviation 0.1. The cameraman image was degraded by blurring with a 25 by 25 11th order spline followed by additive white Gaussian noise of standard deviation 5. In both cases the proposed deconvolution algorithm was compared against iterated soft thresholding using the undecimated wavelet transform. The Daubechies 4 tap wavelet was used to a depth of 3 for the piecewise polynomial image and a depth of 6 for the cameraman image.




Fig. 3: Deconvolution of piecewise polynomial images

4. CONCLUSIONS

We have proposed image restoration algorithms based on a sparse approximation using a quadtree decomposition. Preliminary results suggest that the performance of the denoising algorithm is comparable to soft thresholding of the undecimated wavelet transform on real world images and the deconvolution algorithm can outperform iterated soft thresholding when the variance of the noise is quite large. We are currently investigating adding different tile models to improve both algorithms performance in regions of texture and also generally improving the deconvolution algorithm. We are also interested in the computational complexity of the algorithms and possible speed ups particularly in the deconvolution case.

5. REFERENCES

- [1] I. Daubechies, M. Defrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Comm. Pure Appl. Math.*, vol. 57, no. 11, pp. 1413–1457, 2004.
- [2] M. Figueiredo, J. Bioucas-Dias, and R. Nowak, "Majorization-minimization algorithms for wavelet-based image restoration," *Image Processing, IEEE Transactions on*, vol. 16, no. 12, pp. 2980–2991, 2007.
- [3] M. Figueiredo and R. Nowak, "A bound optimization approach to wavelet-based image deconvolution," *Image Processing,*




Fig. 4: Deconvolution of cameraman image (512×512)

2005. *ICIP 2005. IEEE International Conference on*, vol. 2, pp. II– 782–5, 2005.
- [4] K. Lange, D.R. Hunter, and I. Yang, "Optimization transfer using surrogate objective functions," *J. Comput. Graph. Statist.*, vol. 9, no. 1, pp. 1–59, 2000.
- [5] M. Elad, B. Matalon, J. Shtok, and M. Zibulevsky, "A wide-angle view at iterated shrinkage algorithms," *Wavelets XII. Proceedings of the SPIE*, vol. 6701, pp. 670102, 2007.
- [6] T. Blumensath and M. Davies, "Iterative thresholding for sparse approximations," *Journal of Fourier Analysis and Applications*, Jan 2008.
- [7] R. Shukla, P.L. Dragotti, M.N. Do, and M. Vetterli, "Rate-distortion optimized tree-structured compression algorithms for piecewise polynomial images," *Image Processing, IEEE Transactions on*, vol. 14, no. 3, pp. 343 – 359, Mar 2005.
- [8] R. Shukla and M. Vetterli, "Geometrical image denoising using quadtree segmentation," *Image Processing, 2004. ICIP '04. 2004 International Conference on*, vol. 2, pp. 1213– 1216 Vol.2, 2004.
- [9] A. Scholefield and P.L. Dragotti, "Quadtree structured restoration algorithms for piecewise polynomial images," *to appear in IEEE International Conference on Acoustics, Speech and Signal Processing*, 2009.
- [10] D.R. Hunter and K. Lange, "A tutorial on MM algorithms," *Amer. Statist.*, vol. 58, no. 1, pp. 30–37, 2004.