QUADTREE STRUCTURED RESTORATION ALGORITHMS FOR PIECEWISE POLYNOMIAL IMAGES

Adam Scholefield and Pier Luigi Dragotti

Communications and Signal Processing Group, Electrical and Electronic Engineering,
Imperial College London, Exhibition Road, London SW7 2AZ, England.
Email: adam.scholefield@imperial.ac.uk, p.dragotti@imperial.ac.uk

ABSTRACT

Iterative shrinkage of sparse and redundant representations are at the heart of many state of the art denoising and deconvolution algorithms. They assume the signal is well approximated by a few elements from an overcomplete basis of a linear space. If one instead selects the elements from a nonlinear manifold it is possible to more efficiently represent piecewise polynomial signals. This suggests that image restoration algorithms based around nonlinear transformations could provide better results for this class of signals. This paper uses iterative shrinkage ideas and a nonlinear quadtree decomposition to develop image restoration algorithms suitable for piecewise polynomial images.

Index Terms— Deconvolution, image restoration, piecewise polynomial approximation, quadtrees.

1. INTRODUCTION

Denoising and deconvolution are classic well studied problems that occur in many situations. Commonly one assumes the following degradation model:

$$y = Hx + e, \tag{1}$$

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. In most cases H is ill-conditioned and the pure inverse approximation, $\hat{x} = H^{-1}y$ is heavily corrupted by the coloured noise component $H^{-1}e$. Traditional frequency based techniques such as the Wiener filter overcome this by suppressing $H^{-1}y$ at high frequencies where the inverse is dominated by the coloured noise component. Unfortunately in many applications these high frequency components also contain much of the information of x. This is particularly true in images where edges represent an important part of the visual information.

Because of this the image processing community has moved away from frequency based methods and started using transforms that provide a sparser representation of real world images. The wavelet transform being by far the most common example. Usually an approximation of x is obtained by solving a minimising problem of the form:

$$\hat{\theta} = arg \min_{\theta} \|y - HD\theta\|_2^2 + \lambda\|\theta\|_p, \tag{2}$$

where $x = D\theta$, 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. In the wavelet transform case D is simply the inverse wavelet transform. It is expected that the coefficients vector θ will be sparse.

There has been much recent interest in the solution of (2). In the simple denoising case (i.e. $H = I$) and D a unitary transform the cost function is exactly minimised by simple shrinkage (i.e. hard-thresholding for $p = 0$ and soft-thresholding for $p = 1$). However in the more general case the problem is much more complex. Daubechies et al [1] proved that (2) can be exactly solved for the case where $1 \leq p \leq 2$ by iteratively minimising surrogate cost functions. Similar algorithms have also been derived using for example expectation maximisation (EM) [2], majorisation minimisation (MM), bound optimisation and optimisation transfer algorithms [3, 4, 5]. A good overview of iterated shrinkage is presented in [6]. Blumensath and Davies [7] also show promising results for the $p = 0$ case, however they can no longer guarantee global convergence due to the lack of convexity in the cost function.

Although wavelets have become the standard in image processing applications 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 [8] 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 on real world images. This was due to the sparser representation achieved by their transformation. Willett and Nowak [9] used a piecewise linear wedgelet model to produce an image restoration algorithm for photon-limited medical imaging. Their image deconvolution algorithm uses the iterative

978-1-4244-2354-5/09/$25.00

shrinkage idea derived from an EM approach. In this paper we develop an image restoration algorithm more suitable to standard images using the more complex piecewise polynomial model of [8]. Unlike [8] which uses a rate-distortion cost function more suitable to compression we minimise a different cost function very similar to (2). Shukla et al also developed a denoising algorithm [10] using the piecewise polynomial model however they again used their rate-distortion cost function. The rest of this paper is organised as follows. In Section 2.1 the denoising algorithm is explained, then Section 2.2. develops an iterative deconvolution algorithm using a surrogate function and MM ideas. We also develop a technique to perturbate from local minimum to an approximation with a smaller cost. Section 3 shows the results of this new deconvolution algorithm on piecewise polynomial images and shows comparisons with existing techniques. Finally Section 4 gives conclusions.

2. PROPOSED IMAGE RESTORATION ALGORITHMS

2.1. Quadtree decomposition denoising algorithm for piecewise polynomial images

In nonlinear approximation theory quadtree decomposition is well established as a technique to adaptively partition a two dimensional function. The more flexible nonlinear transformation allows [8] to achieve highly competitive compression and we aim to achieve similar results in the restoration case. In the following we will describe the quadtree decomposition algorithm of [8] with our modifications making it more applicable to denoising. The aim of the algorithm is to solve the following minimisation problem:

$$\hat{\theta} = arg\min_{\theta} \|y - D(\theta)\|_2^2 + \lambda\|\theta\|_0, \qquad (3)$$

where $D(\theta)$ is the image representation with coefficients θ. By $\|\theta\|_0$ we mean the number of non zero coefficients used to produce the image $D(\theta)$. The image partitions shown in figures 1a and 1b are examples of the prune and prune-join models which are explained below.

Algorithm 1 is used to approximately solve (3) using the prune only model resulting in what we will call the pruned quadtree. As figure 1a shows a limitation of the prune only model is that neighbouring regions can only be jointly represented by their parent. It would be beneficial however to allow any neighbouring regions to be jointly represented. To achieve this every leaf of the pruned quadtree is considered to be joined to its neighbouring leaves using algorithm 2 resulting in the partitions of figure 1b.

When approximating a region of the image a hard thresholding step is used, this is optimal as the polynomials we use are orthogonal and the problem is thus linear and unitary for a particular node with a particular boundary.

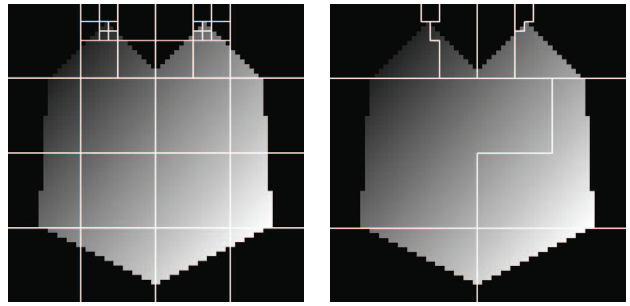

(a) Prune regions (b) Prune-join regions

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

Algorithm 1 (Denoise prune only)

1. *Decompose the image to a predetermined depth.*
2. *Approximate each leaf by two 2-D polynomials separated by a polynomial boundary.*
   (a) *Find the best coefficients for each of a predetermined set of boundaries.*
   (b) *Hard threshold the coefficients of the polynomials for each boundary.*
   (c) *Choose the boundary that results in the minimum cost as described by* (3).
3. *Approximate the parents of the leaves using the same approach as 2.*
4. *If the parent has a smaller cost than the sum of the costs of the four children then prune the children and use the parents approximation.*
5. *Repeat step 3 and 4 all the way up the tree.*

Algorithm 2 (Denoise join)

1. *Visit the leaves of the pruned quadtree in a top to bottom left to right manner.*
2. *Join a leaf with a neighbouring leaf later in the tree if the joined cost is less than the sum of the individual costs of the two leaves .*
3. *If two leaves are joined the two leaves are considered as one leaf for the rest of the joining process.*

2.2. Deconvolution using the non linear quadtree decomposition transformation

To extend the quadtree decomposition algorithm for deconvolution we attempt to minimise a cost function virtually identical to (2) but with slight modification for the non linear case:

$$\hat{\theta} = arg\min_{\theta} \|y - HD(\theta)\|_2^2 + \lambda\|\theta\|_0. \qquad (4)$$

The H causes all the basis functions in our transformation to overlap which means that we cannot locally look for the best tile. This is equivalent to the non-unitary linear denoising problem where all the equations are coupled together. As the linear case we use a surrogate function and the MM philosophy to decouple these equations, for a good introduction to MM algorithms see [11]. Equations (5) and (6) show the original cost function C, and surrogate cost function C_{sur} respectively.

$$C(D(\theta)) = \|y - HD(\theta)\|_2^2 + \lambda\|\theta\|_0 \tag{5}$$

$$\begin{aligned} C_{sur}(D(\theta) \mid a) = C(D(\theta)) - \|HD(\theta) - Ha\|_2^2 \\ + \alpha\|D(\theta) - a\|_2^2 \end{aligned} \tag{6}$$

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

$$C_{sur}(D(\theta) \mid a) \geq C(D(\theta)) \qquad \forall\theta \tag{7}$$

$$C_{sur}(a \mid a) = C(a) \tag{8}$$

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

$$\begin{aligned} C_{sur}(D(\theta) \mid a) = \|y\|_2^2 - 2D(\theta)^T H^T y + \|HD(\theta)\|_2^2 \\ - \|HD(\theta)\|_2^2 + 2D(\theta)^T H^T Ha \\ - \|Ha\|_2^2 + \alpha\|D(\theta)\|_2^2 + \alpha\|a\|_2^2 \\ - 2\alpha D(\theta)^T a + \lambda\|\theta\|_0 \end{aligned} \tag{9}$$

$$\begin{aligned} \frac{C_{sur}(D(\theta) \mid a)}{\alpha} = \left\| a + \frac{H^T}{\alpha}(y - Ha) - D(\theta) \right\|_2^2 + \bar{\lambda}\|\theta\|_0 \\ + \text{terms independent of } \theta \end{aligned} \tag{10}$$

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:

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

From the inequalities of a maximiser we know that

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

$$C_{sur}(D(\theta^i) \mid D(\theta^i)) = C(D(\theta^i)) \tag{13}$$

So if the denoising algorithm guarantees that $C_{sur}(D(\theta^{i+1}) \mid D(\theta^i)) \leq C_{sur}(D(\theta^i) \mid D(\theta^i))$ then the sequence is guaranteed to be decreasing. As the previously introduced denoising algorithm only approximately solves (3) then this is not guaranteed in its current state. To achieve convergence we use an update algorithm which starts looking for the best approximation from the current representation θ^i. This not only guarantees a decreasing sequence but it is also more computationally efficient.

As the denoising case we first assume the simpler prune only model and only introducing joining when our deconvolution pruning algorithm (algorithm 3) cannot further improve the representation.

Algorithm 3 (Deconvolution, prune only model)

1. *Denoise pruning only. (Algorithm 1).*
2. *Iterate linear hard thresholding.*
   (a) *Find basis functions for the current tile and boundary structure.*
   (b) *Find the coefficients for updated image $D(\theta^i) + \frac{H^T}{\alpha}(y - HD(\theta^i))$ using these basis functions.*
   (c) *Hard threshold the coefficients.*
   (d) *If the coefficients have changed goto 2b.*
3. *Update pruned quadtree.*
   (a) *Inspect each leaf in a bottom up approach.*
   (b) *Find the best boundary model for the leaf.*
   (c) *Split the leaf to its four children if the sum of the children's cost is less than that of the current leaf.*
   (d) *If four leaves have the same parent and the parent has a smaller cost than the sum of its children's costs then prune the leaves.*
   (e) *If after doing 3a-d for the whole tree a boundary or tile structure has changed then goto step 2.*
4. *Update a single tile or prune four children using* (5) *directly.*
   (a) *Visit every leaf in the tree.*
   (b) *Remove the blurred tile of this leaf from the current representation and find the residual over the region that the blurred tile covers.*
   (c) *Find the best blurred tile that approximates the residual by trying every boundary and hard thresholding as before.*
   (d) *If four leaves have the same parent then remove the four leaves and find the best blurred parent to replace them as steps 4b,c.*
   (e) *After doing steps 4a-d for the whole tree update the representation that decreases the cost the most and goto step 2. If no improvement can be made then terminate.*

The iterate linear hard thresholding step of algorithm 3 only calculates the basis functions once and it can thus very quickly update the coefficients until convergence. This allows step 3 to just look to improve the boundaries and tile structure which is a slower process. When step 4 is reached the surrogate approach has reached a local minimum so we attempt to improve further by using (5) directly. Although all the basis functions in (5) overlap preventing them all to be solved in a reasonable amount of time it is still possible to update only one and fix all the others. It is this principle that allows step 4 to update a single tile or replace four children with their parent. Only the single representation that decreases the cost the most is updated.

The algorithm to join the pruned quadtree is very similar to algorithm 3. First the pruned quadtree is joined and then hard thresholding is iterated with the tiles and boundaries fixed. A slower algorithm is then used to improve the

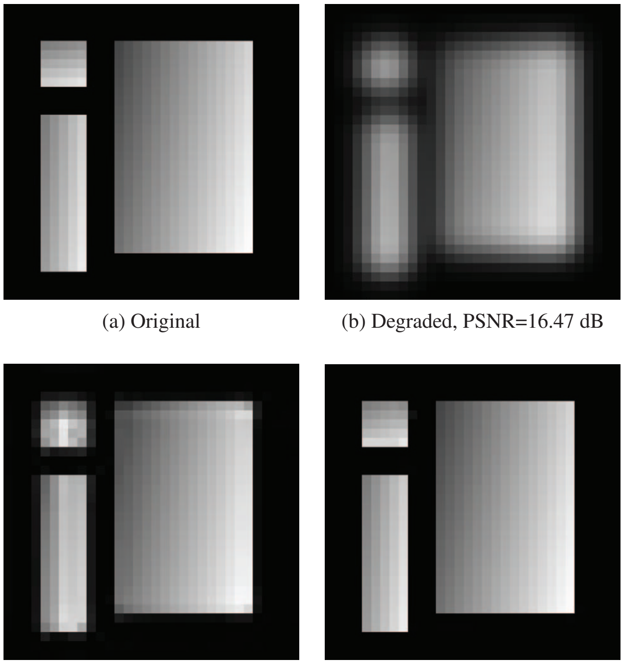

(a) Original

(b) Degraded, PSNR=16.47 dB

(c) Reconstructed with wavelets, PSNR=27.38 dB

(d) Reconstructed with proposed algorithm, PSNR=50.22 dB

Fig. 2: Deconvolution of piecewise polynomial images

boundaries and tile structure using the surrogate approach and finally a single tile or two neighbouring tiles are replaced by a single tile to escape from local minimum. The algorithms are iterated in the same way as algorithm 3.

3. EXPERIMENTAL RESULTS

The proposed algorithm can be implemented to any polynomial degree with any set of continuous boundaries however in the following results we use polynomials of maximum degree 1 and straight edge boundaries. Figure 2 shows the deconvolution of a synthetic piecewise linear image that has been blurred by a 7 by 7 quadratic spline followed by additive Gaussian noise with a standard deviation of 0.1. The proposed algorithm is compared against iterated soft thresholding using the stationary wavelet transform. The Daubechies 4 tap wavelet was used to a depth of 3.

4. CONCLUSIONS

We have presented image restoration algorithms based on iterative shrinkage ideas and a nonlinear quadtree decomposition for use on piecewise polynomial images. Preliminary results suggest the success of the algorithm for reconstructing this class of signals and we are currently investigating extending these ideas to real world images.

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 and R. Nowak, "An EM algorithm for wavelet-based image restoration," *Image Processing, IEEE Transactions on*, vol. 12, no. 8, pp. 906– 916, 2003.

[3] 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.

[4] M. Figueiredo and R. Nowak, "A bound optimization approach to wavelet-based image deconvolution," *Image Processing, 2005. ICIP 2005. IEEE International Conference on*, vol. 2, pp. II– 782–5, 2005.

[5] 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.

[6] 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.

[7] T. Blumensath and M. Davies, "Iterative thresholding for sparse approximations," *Journal of Fourier Analysis and Applications*, Jan 2008.

[8] 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.

[9] R.M. Willett and R. Nowak, "Platelets: a multiscale approach for recovering edges and surfaces in photon-limited medical imaging," *Medical Imaging, IEEE Transactions on*, vol. 22, no. 3, pp. 332– 350, 2003.

[10] 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.

[11] D.R. Hunter and K. Lange, "A tutorial on MM algorithms," *Amer. Statist.*, vol. 58, no. 1, pp. 30–37, 2004.