

Multigrid Adaptive Image Processing

Michael Unser

Biomedical Engineering and Instrumentation Program, NCRP, Bldg. 13, Room 3N17
National Institutes of Health, Bethesda, MD 20892-5766 USA

ABSTRACT

We consider a general weighted least squares approximation problem with a membrane spline regularization term. The key parameters in this formulation are the weighting factors which provide the possibility of a spatial adaptation. We prove that the corresponding space-varying variational problem is well posed, and propose a novel multigrid computational solution. This multiresolution relaxation scheme uses three image pyramids (input data, weights, and current solution) and allows for a very efficient computation with an effective $O(N)$ complexity, where N is the number of pixels. This general multigrid solver can be useful for a variety of image processing tasks. In particular, we propose new multigrid solutions for noise reduction in images (adaptive smoothing spline), interpolation/reconstruction of missing image data, and image segmentation using an adaptive extension of the K -means clustering algorithm.

1. INTRODUCTION

Variational methods and the use of optimization principles in general provide attractive tools for deriving many image processing algorithms [1, 2]. The main idea is to convert the initial task into an energy (or functional) minimization problem, which presents certain conceptual advantages. Typical energy functions include two components: a data term that forces the solution to be consistent with the available measurements, and a regularization term that imposes some smoothness on the solution. The second term is especially important for obtaining a problem that is not ill-posed mathematically [3]. In a Bayesian framework, it is determined by our a priori knowledge of the solution [4].

Unfortunately, even in the simplest case of quadratic energy functions (linear solution), the minimization of these functionals tends to be computationally quite demanding. The matrices involved are too large to be handled directly, and the solutions are computed iteratively using standard numerical methods (e.g., Jacobi, SOR, Conjugate gradient). These methods typically require a large number of iterations, which

limits their applicability even though convergence is usually guaranteed. One important exception is the space-invariant case where the solution can be determined directly by digital filtering using the FFT ($O(N \log N)$) [5]. Recursive filtering solutions ($O(N)$) have also been proposed for the simplest forms of energies [6].

As a rule, we can gain in computational efficiency only if we take advantage of certain special features of the problem at hand, e.g., space-invariance. Of special interest is the case when the functional involves simple differential operators so that the solution can be specified by a partial differential equation, as is often the case in computer vision [7]. This type of problem can be solved very efficiently by using some of the recent multigrid relaxation schemes developed in applied mathematics [8, 9]. Terzopoulos has demonstrated the usefulness of this approach in a variety of low level computer vision problems: visual surface reconstruction, lightness, shape-from-shading, and optical flow [10, 11].

Our intent in this paper is to show that the multigrid methodology can also be beneficial for some of the more traditional image processing tasks (interpolation, noise reduction, and segmentation). The key motivation is to move away from the standard space-invariant type of processing, without necessarily having to suffer from a corresponding increase in computational complexity. We are also interested in the variational approach because of its generality and its conceptual appeal, namely, the fact that the design process boils down to the specification of an energy functional that is appropriate for the application at hand.

For this purpose, we will consider the problem of the minimization of a general criterion of the form:

$$E(u) = \sum_{(k,l) \in Z^2} w(k,l) (f(k,l) - u(k,l))^2 + \lambda \sum_{(k,l) \in Z^2} (d_x * u(k,l))^2 + (d_y * u(k,l))^2, \quad (1)$$

where f is the input image, u is the desired solution, and $w \geq 0$ is a map of space-varying weights; d_x and d_y are the horizontal

and vertical first difference convolution operators, respectively. The second space-invariant term in (1) is a membrane spline regularizer; the amount of smoothness is controlled by the parameter λ . The versatility of the criterion comes from the fact that the weights w can vary adaptively over the image; they may, for example, represent a mask that emphasizes certain regions (or points) of interest. In Section 2, we will first show that the problem is well posed mathematically. We will then describe our multigrid computational solution. In Section 3, we will be more specific and show how various image processing tasks can be cast into this general framework. We will propose several examples of multigrid adaptive image processing, most of which we believe are original.

2. GENERAL FORMULATION

2.1 Mathematical analysis

By taking the partial derivative of (1) with respect to u , we find that u is the solution of the difference equation

$$f_w(k,l) = w(k,l) \cdot u(k,l) + \lambda \cdot L * u(k,l), \quad (2)$$

where $f_w(k,l) = w(k,l) \cdot f(k,l)$ and where $L = d_x^T * d_x + d_y^T * d_y$ is the discrete Laplacian operator. Note that the homogenous part of this equation is similar to Helmholtz's reduced wave equation $k^2 u + Lu = 0$, except for the weighting factor that is space-varying. The equivalent form of (2) in matrix notation is

$$f_w = Wu + \lambda Lu = Au \quad (3)$$

where W is the diagonal weight matrix, $f_w = Wf$ the weighted data vector, and $A = W + \lambda L$ a symmetrical positive definite matrix. The positivity constraint on W and the special form of the regularization functional L guarantees that A is strictly diagonal dominant, and therefore non-singular. This condition insures that our minimization problem is well posed, and that it has a unique solution :

$$u = A^{-1} f_w. \quad (4)$$

The corresponding minimum value of the energy criterion is

$$\min_{\mathbf{v}} E(\mathbf{v}) = E(\mathbf{u}) = \mathbf{f}^T W \mathbf{f} - \mathbf{f}^T W \mathbf{u} = \mathbf{f}_w^T (\mathbf{f} - \mathbf{u}). \quad (5)$$

2.2 Multigrid Implementation

Even if the problem is well posed mathematically, the size of the matrices is such that the linear system of equations (3) cannot be inverted directly. To derive an efficient numerical procedure, we take advantage of the fact that (2) is scalable and therefore suitable for multigrid processing. For this purpose, we define three image pyramids : one for the forcing term f_w or the current residual $r = f_w - wu + \lambda L * u$, one for the spatial weights w , and one for the computation of the solution u . The w pyramid is specific to our formulation and is generated during

the initialization phase by successive filtering and decimation by a factor of two in each direction (restriction or DOWN operation) . The two other pyramids u and r are computed and updated during the course of iteration : the current solution update u may either be interpolated to the next finer level by upsampling and filtering (UP operation), or its residual r projected onto the next coarser level using the DOWN operation. In either case, we used piecewise linear filters (first degree splines) which are consistent with the order of the differential operator. It is important to note that the basic idea of multigrid processing is to iterate on a residual system of equations, which provides the appropriate solution update, rather than on the original system itself [9]. At the i th level of the pyramid where the sampling step is $h = 2^i$, the rescaled version of the equation that needs to be solved is

$$r^{(h)} = w^{(h)} u^{(h)} + h^{-2} \lambda (L * u^{(h)}), \quad (5)$$

where the superscript h refers to the data representations at that particular resolution, and where the factor h^{-2} provides the proper scale normalization for the 2nd order differential operator. Instead of solving (3) directly, we use a standard multigrid iteration strategy [8, 9]. Specifically, we apply n_1 iterations of the dampened Jacobi method, which is good for getting the high frequency part of the solution and is also guaranteed to converge in our particular case. The residual is then projected onto the next coarser level and the process is iterated. At the bottom of the pyramid, the equation is solved within the required level of precision. This coarse part of the solution is then interpolated and used to update the intermediate finer step solutions, with n_2 additional Jacobi iterations per level.

For most cases that we tested, we found that the equation could be solved within an acceptable error tolerance with only one full multigrid V-cycle with $n_1 = 1$ and $n_2 = 2$. Since the single-scale version of the Jacobi algorithm typically requires hundreds of iterations, the present approach provides orders of magnitude speed improvement (one full multigrid V-cycle is comparable in complexity to $2n_1 + 2n_2$ fine-grid iterations). In essence, we have an $O(N)$ algorithm that is extremely competitive.

3. IMAGE PROCESSING EXAMPLES

Now that we have a general multigrid solver of (2) at our disposal, we will look at some specific examples.

3.1 Smoothing spline filter

By setting $w=1$ over the whole image, we get the equivalent of a non-separable smoothing spline filter [6, 12].

The corresponding space-invariant operator can be characterized by its 2D transfer function

$$H(z_1, z_2) = \frac{1}{1 + \lambda(-z_1 + 2 - z_1^{-1} - z_2 + 2 - z_2^{-1})}. \quad (6)$$

This is an exponential filter with a response that is approximately isotropic. The amount of smoothing is directly controlled by the parameter λ . Specifically, we can show that the filter's equivalent window size (standard deviation) is

$$\sigma_\lambda = \sqrt{2\lambda}, \quad (7)$$

a relation that gives us a quantitative understanding of the smoothing effect of λ . Depending on the application, the optimal regularization parameter may also be determined from the data using cross-validation methods [13], or from a given measurement model (signal + noise) [14]. For the case in which the image ($f=s+n$) is corrupted by additive white noise with a known variance σ^2 , we have derived a simple quasi-Wiener rule for the selection of the optimum regularization parameter

$$\lambda = \frac{\sigma^2}{E\{f \cdot Lf\} - 4\sigma^2}, \quad (8)$$

where $E\{f \cdot Lf\}$ denotes an estimate of the correlation between the noisy image data and its Laplacian as defined in (2). For the images that we have considered so far, this simple rule works as well as cross-validation, and is certainly much less time consuming.

An interesting extension is to use space-varying regularization factors $\lambda(k, l)$ that are adjusted locally using signal statistics evaluated over a sliding window [14]. Note that this adaptive formulation also fits into the present framework: the spatial weights in (2) are simply $w(k, l) = \lambda_0 / \lambda(k, l)$, with the constraint that $\lambda(k, l) > 0$. We are currently considering this approach for designing an adaptive Wiener filter for additive white noise which may be useful for processing noisy electron micrographs.

In our current implementation, such a smoothing spline filtering of a 256×256 image (one full multigrid V -cycle) is approximately twice as fast as a comparable Fourier domain filtering with a proper handling of boundary conditions (512×512 FFTs). In addition, our algorithm works for the space-variant case as well, and its computational efficiency is unaffected.

3.2 Interpolation and image reconstruction

Equation (1) is also suitable for surface reconstruction and extrapolation. In this case, the weights are set to zero where no data is available and one otherwise. The value of λ then determines the tightness of the fit at the data points, while the

surface is extrapolated such that its Laplacian is zero elsewhere. Fig. 1 presents an example of the reconstruction of an image from the values of its contour points only. In this example, the edge points were selected by simple thresholding of the image gradient. We believe that this type of edge-based image representation can be further improved by using a detection method that is more directly related to the reconstruction algorithm, and also by closing some of the contour lines to avoid blending artifacts.

We have also used the same reconstruction technique to interpolate non-uniformly and sparsely sampled data points, with very satisfactory results.

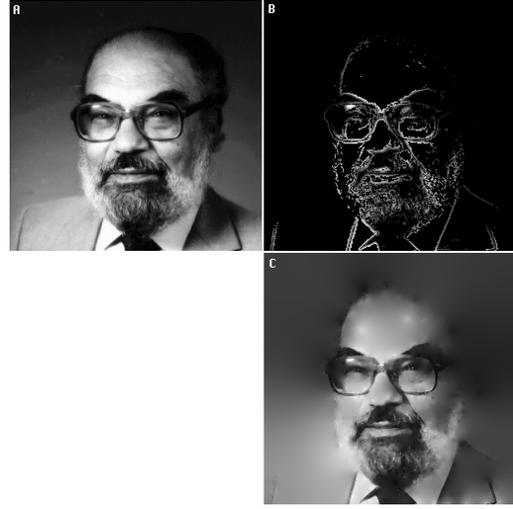


Fig. 1 : (a) Initial 256×256 image, (b) Edge map: non-edge points are simply set to zero, (c) reconstructed image using the proposed algorithm. Non-edge points are extrapolated such that their Laplacian is zero.

3.3 Adaptive K -means image segmentation

The last application that we present is an adaptive K -means segmentation algorithm similar to the method described by Pappas [15]. The essential difference with this earlier work is that we use a membrane model to specify the smoothness of the means over the regions rather than a sliding window approach. The advantage is mostly conceptual because the model is now entirely specified through a single Gibbs energy function. Assuming that there are K image regions R_1, R_2, \dots, R_K with slowly varying means u_1, u_2, \dots, u_K , we want to find the partition that minimizes

$$E(u_1, u_2, \dots, u_K) = \sum_{i=1}^K \left(\sum_{(k,l) \in R_i} [f(k,l) - u_i(k,l)]^2 + \lambda_i \sum_{(k,l) \in \mathbb{Z}^2} (d_x * u_i(k,l))^2 + (d_y * u_i(k,l))^2 \right). \quad (9)$$

The solution is determined iteratively by successively optimizing over R_1, R_2, \dots, R_K keeping the u_i 's fixed (nearest neighbor rule), and then over u_1, u_2, \dots, u_K keeping the R_i 's fixed. The second step is solved using the algorithm described above over each region with $w_i(k,l)=1$ when $(k,l) \in R_i$ and $w_i(k,l)=0$ otherwise. Note that with this formulation the mean signals are defined everywhere. We observe that this procedure is equivalent to the standard K -means clustering when the λ_i 's are sufficiently large, in which case u_i converges to a constant: the mean value over the corresponding region. Typically, we use the standard K -means segmentation to provide a starting point for the algorithm.

An example of adaptive image segmentation is shown in Fig. 2d, next to the conventional K -means solution in Fig. 2b. The adaptive scheme does a better job in preserving the local image details which are important for the interpretation. The corresponding slowly-varying region model is displayed in Fig. 2c.

The algorithm can also be made robust with respect to noise by including an additional Gibbs spatial interaction term in the energy (clique potentials), as described in [15]. The additional clique parameters will essentially provide a controlling mechanism for the removal of isolated points. This can be achieved by a simple modification of the region assignment part of the algorithm, using the iterated conditional modes approach proposed by Besag [16].

4. CONCLUSION

Our objective was to demonstrate some of the potential uses and advantages of multigrid methods for image processing. For this purpose, we started by developing a multigrid solver for a space-varying form of Helmholtz's equation. Our primary motivation was to investigate the possibility of local adaptation in the solution of regularization problems, and to move away from the standard space-invariant type of processing. This approach turned out to be quite fruitful. Our preliminary results suggest that the proposed methodology should be useful for a variety of image processing problems.

The multigrid algorithm that we propose here is computationally very efficient. For all practical purposes, it is an $O(N)$ method where N is the number of pixels in the image. In fact, in the space-invariant case, the approach is even competitive with Fourier filtering techniques (e.g. smoothing spline filter).

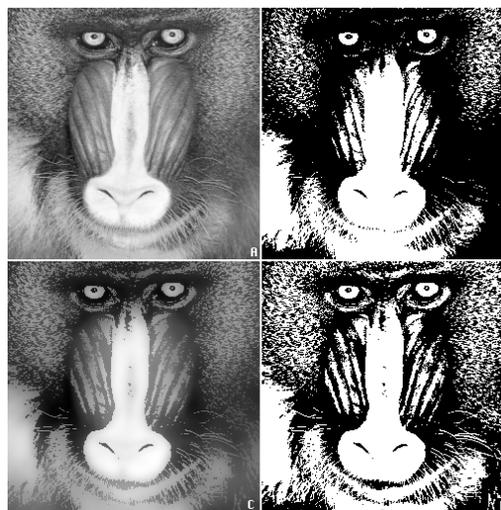


Fig. 2 : (a) 256x256 test image, (b) standard K -means segmentation ($K=2$), (c) fitted image model with $\lambda=16$; (d) adaptive K -means segmentation.

References

- [1] T. Poggio, V. Torre and C. Koch, "Computational vision and regularization theory," *Nature*, vol. 317, pp. 314-319, 1985.
- [2] A. Blake and A. Zisserman, *Visual reconstruction*. Boston, MA: MIT Press, 1987.
- [3] A.N. Tikhonov and V.Y. Arsenin, *Solution of ill-posed problems*. Washington, DC: Winston and Sons, 1977.
- [4] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images," *IEEE Trans. Patt. Anal. Machine Intell.*, vol. PAMI-6, no. 6, pp. 712-741, November 1984.
- [5] W.K. Pratt, *Digital image processing*. New York: Wiley, 1978.
- [6] M. Unser, A. Aldroubi and M. Eden, "Recursive regularization filters: design, properties and applications," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 3, pp. 272-277, March 1991.
- [7] B.K.P. Horn, *Robot Vision*. New York: McGraw-Hill, 1986.
- [8] W. Hackbush, *Multi-Grid Methods and Applications*. New York: Springer-Verlag, 1985.
- [9] W.L. Briggs, *A multigrid tutorial*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1987.
- [10] D. Terzopoulos, "Multilevel computational processes for visual surface reconstruction," *Comput. Vision, Graphics, Image Processing*, vol. 24, pp. 52-96, 1983.
- [11] D. Terzopoulos, "Image analysis using multigrid relaxation methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-8, no. 2, pp. 129-139, March 1986.
- [12] T. Poggio, H. Voorhees and A. Yuille, "A regularized solution to edge detection," *Journal of Complexity*, vol. 4, no. 2, pp. 106-123, 1988.
- [13] G. Wahba, "Practical approximate solutions to linear operator equation when the data are noisy," *SIAM J. Math. Anal.*, vol. 14, no. 4, pp. 651-667, 1977.
- [14] S.J. Reeves, "Optimal space-varying regularization in iterative image restoration," *IEEE Trans. Image Processing*, vol. 3, no. 3, pp. 319-324, May 1994.
- [15] T.N. Pappas, "An adaptive clustering algorithm for image segmentation," *IEEE Trans. Signal Processing*, vol. 40, no. 4, pp. 901-914, April 1992.
- [16] J. Besag, "On the statistical analysis of dirty pictures," *J. Roy. Stat. Soc. B*, vol. 48, no. 3, pp. 256-302, 1986.