

MULTIGRID IMAGE RECONSTRUCTION FROM ARBITRARILY SPACED SAMPLES

Muthuvel Arigovindan, Michael Sühling, Patrick Hunziker, and Michael Unser*

Swiss Federal Institute of Technology Lausanne
Biomedical Imaging Group, IOA
CH-1015 Lausanne, Switzerland

e-mail: muthuvel.arigovindan@epfl.ch, michael.suehling@epfl.ch, michael.unser@epfl.ch

ABSTRACT

We propose a novel multiresolution-multigrid based signal reconstruction method from arbitrarily spaced samples. The signal is reconstructed on a uniform grid using B-splines basis functions. The computation of spline weights is formulated as a variational problem. Specifically, we minimize a cost that is a weighted sum of two terms: (i) the sum of squared errors at the specified points; (ii) a quadratic functional that penalizes the lack of smoothness. The problem is equivalent to solving a very large system of linear equations, with the dimension equal to the number of grid points. We develop a computationally efficient multiresolution-multigrid scheme for solving the system. We demonstrate the method with image reconstruction from contour points.

1. INTRODUCTION

The problem of signal reconstruction from irregularly spaced samples occurs frequently in signal and image processing. Examples include shape reconstruction, landmark interpolation, image recovery from the contours etc. For our part, we are motivated by the problem of reconstructing 3D echocardiograms from series of 2D hand-held acquisitions spanning the volume of interest [1].

The difficulty with this kind of reconstruction, especially when the data is sparse, is that the problem is ill-posed. Therefore, to get a meaningful reconstruction, one usually needs to introduce some regularization constraint. An elegant analytical solution is provided by the method of thin plate splines which produces a continuously-defined solution that minimizes a Laplacian-related semi-norm [2]. In other words, the problem is set up such as to find a compromise between fitting the data well and penalizing reconstructions that are not smooth enough. The solution is expressed as a linear combination of shifted radial basis functions positioned at the location of the data points [2, 3, 4].

*P. Hunziker is with the Kantonsspital Basel, Petersgraben 5, CH-4031 Basel, e-mail: PHunziker@uhbs.ch.
This work was supported in part by the Swiss Heart Foundation and the Swiss National Science Foundation under grant 3200-059517.99.

The optimal weights are determined as solution of a linear system of equations.

While thin plate splines and radial basis functions (RBF) are undoubtedly among the preferred techniques for dealing with non-uniform interpolation in multiple dimensions [5, 6], they tend to break down numerically when there are too many data points; say, $M > 100$. The main difficulty is that the linear system to be solved is ill-conditioned. Moreover, the matrix is not sparse. Hence, solving the system soon becomes overly expensive; the complexity is in $\mathcal{O}(M^3)$, where M is the number of sample points. Various solutions have been proposed for reducing this complexity and improving the numerical behavior [7], but there is still a long way to go for making them practical. Another fundamental limitation is that computing the weights is only part of the effort. Indeed, the solution has to be resampled numerically if it is to be displayed on a regular grid; this will cost an additional $\mathcal{O}(MN)$, where N is the number of grid points.

In 1D, the situation is not as dramatic because the optimal solution can be expressed in terms of non-uniform B-splines which are compactly supported as opposed to the RBFs which are not; this is the key for obtaining an efficient solution [8]. Unfortunately, this is not possible in higher dimensions; i.e., there are no compactly supported functions that reproduce spline-related RBFs.

In this paper, we circumvent the above mentioned difficulties by searching for the solution of the regularized reconstruction problem in the space of uniform spline where the basis functions are now attached to the reconstruction grid, as opposed to the data points. In other words, we are proposing to discretize thin plate splines using uniform B-splines with a degree that is matched to the underlying semi-norm. In this way, we have at least the guarantee that the solution in 1D coincides with the theoretical one, provided that the sample points are on the reconstruction grid. This helps to say qualitatively that the solution in general will not be very different from the exact analytical one.

The proposed approach has many advantages over using RBFs. First, the linear system for getting the B-spline

coefficients is well-conditioned. Second, the system matrix is sparse resulting in much faster computation. Third, the formulation lends itself quite naturally to an efficient multiresolution and multigrid implementation, as shown in Section 3. This reduces the complexity of solving the linear system to $\mathcal{O}(N)$, where N is the number of grid points. Fourth, after solving the linear system, there is no expensive resampling step as in RBF reconstruction. With linear B-splines, the resampling is free, while for cubic splines, it can be achieved by simple FIR post-filtering [9]. Finally, we should note that unlike other purely discrete variational formulations of the image reconstruction problem [10, 11, 12], we get a solution that is defined in the continuous domain with all the corresponding advantages.

The present algorithm offers a lot of a user-flexibility. The number and location of the data points can be arbitrary without any incidence on computational speed. One can reduce measurement noise by imposing more or less smoothness on the solution (regularization). One can also reconstruct the image at any desired resolution (step h); the solution will converge to the analytical one (RBF) as h gets sufficiently small—the rate is given by the order of the spline.

2. PROBLEM FORMULATION AND SOLUTION

The method described here is applicable to any number of dimension. For notational convenience, we present it for 2D. Given a set of noisy measurements $\{f_i\}$ of the image at sampling locations $\{x_i, y_i\}$, the problem is to find a uniform spline $S(x, y)$ of the form

$$S(x, y) = \sum_{l=0}^{N_y-1} \sum_{k=0}^{N_x-1} c_{k,l} \beta^n(x/h_x - k) \beta^n(y/h_y - l) \quad (1)$$

such that

$$\sum_i |S(x_i, y_i) - f_i|^2 + \lambda \mathcal{M}(S(x, y)) \quad (2)$$

is minimized. Here $\beta^n(x)$ is the B-spline of degree n [9], and $\mathcal{M}(\cdot)$ is a regularization functional that penalizes non-smooth solutions; typically, the L_2 -norm of some derivatives (gradient or Laplacian). The weighting factor λ allows a trade-off between the smoothness and the reconstruction error at the specified points.

We now introduce the corresponding matrix formulation for an image reconstructed on an $N_x \times N_y$ grid. The goal is to search for the coefficient vector

$$\mathbf{c} = [c_{0,0} \cdots c_{N_x-1,0} \cdots c_{N_x-1,N_y-1}]^T$$

such that (2) is minimized. Defining $B_{N_x l+k}(x, y) = \beta^n(x/h_x - k) \beta^n(y/h_y - l)$, the solution

$$\mathbf{c} = [c_{0,0} \cdots c_{N_x-1,0} \cdots c_{N_x-1,N_y-1}]^T$$

is given by,

$$[\mathbf{S} + \lambda \mathbf{R}_P] \mathbf{c} = \mathbf{b} \quad (3)$$

where

$$\{S\}_{j,k} = \sum_i B_j(x_i, y_i) B_k(x_i, y_i),$$

$$\{b\}_j = \sum_i B_j(x_i, y_i) f_i,$$

and where \mathbf{R} is a circulant matrix with a 2D filter P (mask) that depends on the type of regularization.

For the simplest nontrivial case—i.e. for the linear splines,—the matrix \mathbf{S} is block-band diagonal of width three, and each block is band-diagonal with width three. In this case it is appropriate to use the following regularization:

$$\mathcal{M}_1(S(x, y)) = \int \int \left[\left| \frac{\partial S(x, y)}{\partial x} \right|^2 + \left| \frac{\partial S(x, y)}{\partial y} \right|^2 \right] dx dy$$

The regularization filter turns out to be

$$P(z_1, z_2) = - \left(\frac{h_y}{h_x} \right) B^3(z_1) {}^2(z_2) - \left(\frac{h_x}{h_y} \right) {}^2(z_1) B^3(z_2)$$

where (z) is the first-order difference operator and $B^n(z) = \sum_k \beta^n(k) z^{-k}$. For cubic splines, the matrix \mathbf{S} is block-band diagonal of width seven, and each block is band-diagonal with width seven. In this case, it is appropriate to use

$$\mathcal{M}_2(S(x, y)) = \int \int \left(\frac{\partial^2 S(x, y)}{\partial x^2} + \frac{\partial^2 S(x, y)}{\partial y^2} \right)^2 dx dy$$

as regularization. The regularization filter becomes

$$\begin{aligned} P(z_1, z_2) &= \left(\frac{h_y}{h_x^3} \right) B^7(z_1) B^3(z_2) {}^4(z_2) \\ &+ \left(\frac{1}{h_x h_y} \right) {}^2(z_1) B^5(z_1) {}^2(z_2) B^5(z_2) \\ &+ \left(\frac{h_x}{h_y^3} \right) B^7(z_2) \beta^3(z_1) {}^4(z_1) \end{aligned}$$

3. RECONSTRUCTION ALGORITHM

3.1. Multiresolution Strategy

The system of linear equations (3) is typically very large but sparse. Since the dimension of the system depends upon the resolution of reconstruction, one naturally thinks of multiresolution.

Let the width of the reconstructed image along x and y be equal to L , and let the finest grid spacings be $h_x = h_y = h = L/2^n$. Then $N_x = N_y = N = 2^n + 1$. At any coarser resolution, the grid spacing will be $2^i h$ with $0 < i < n$. Let

$$\mathbf{A}_i \mathbf{c}_i = \mathbf{b}_i \quad (4)$$

be the linear system to be solved at resolution i . The idea is to solve c_{n-1} exactly and to interpolate the solution to the next finer scale using the two-scale relation for B-spline [9]. Specifically, the refinement filters for the linear and cubic splines are $U_1(z) = (z + 2 + z^{-1})/2$ and $U_3(z) = [U_1(z)]^2$, respectively. This serves as initialization for the computation of c_{n-2} and so on. Eventually, c_0 gives the reconstruction at the required resolution.

For this multiresolution strategy, the matrix $A_i = S_i + \lambda R_{P_i}$ has to be computed for all the resolutions. At first sight, it appears that the sample matrix S_i has to be computed for each resolution level i , going back to the sample points each time. This suggests a computational cost for S_i that is proportional to the number of sample points irrespective of the reconstruction grid size. This obviously not good news when number of sample points is large.

Because of the special structure of the problem and an appropriate selection of basis functions, we can prove that the matrices A_i and the vectors b_i are related as follows.

$$A_{i+1} = U_i^T A_i U_i \quad (5)$$

$$b_{i+1} = U_i^T b_i \quad (6)$$

where U_i is a $(2^{n-i} + 1) \times (2^{n-i-1} + 1)$ matrix that corresponds to upsampling and filtering with the binomial filter $U_1(z)$ or $U_3(z)$, two-scale filter. Hence, what we need to compute from the sample points is only the matrix at the finest resolution A_0 . Thanks to this strategy, the computation of all the matrices becomes quite affordable.

3.2. Multigrid Iteration

In the described multiresolution, the solution at each resolution level is obtained iteratively by refinement of the initialization from the coarser level. The refinement can be done using iterative techniques such as Jacobi or Gauss-Seidel. However, when the dimension is large, convergence of such iterators may not be sufficient, even with initialization from the coarser resolution, and hence the reconstruction cannot be completed within affordable time. Therefore we adopt a multigrid strategy [13] as described below.

Multigrid iteration is obtained by using classical iterators as building blocks. We will first describe the classical iterative scheme and then the multigrid iteration. Let $c_i^{(k)}$ be the initialization for the resolution level i . Then it is refined as follows:

$$c_i^{(k+1)} = c_i^{(k)} + \omega \hat{A}_i^{-1} (b_i - A_i c_i^{(k)}) \quad (7)$$

where \hat{A}_i^{-1} denotes the approximate inverse and ω is a damping factor. If \hat{A}_i is the diagonal of A_i , the iterator is called damped-Jacobi and when it is the lower triangular of A_i , the iterator is called damped Gauss-Seidel. Both implementations have same computational complexity. See [13] for a comprehensive treatment of such iterators.

It is important to be note here that such iterators have a smoothing effect on the error. Let $e_i^{(k)}$ be the error after the k th iteration. The larger k , the smoother the error, and for sufficiently large k , there will not be significant improvement in the error anymore. Since the error is smooth after a few relaxations, it can be represented at lower dimension; i.e. on a coarser grid. The advantage is twofold. First, it is computationally efficient to iterate in lower dimension. Second, a smooth error becomes somewhat bumpier at a coarser resolution. Hence, relaxation on the coarser grid further reduces the error. Afterwards one can interpolate the result and use it for correction on the finer grid. This reduces the error significantly. At the same time, in most of the cases, there will be some amount of non-smooth error introduced in the correction procedure. Hence it is customary to do a few relaxations after coarse-grid correction. The overall scheme is computationally more effective than relaxation at the finest grid. However, when the dimension of the system at the coarser grid is large, relaxation will also stall there. Since the problem here is exactly the same as the original one, one can think of applying the same three-step procedure recursively. See [14, 13]. It can be proven that since the system is positive definite and because of the relation (5), this recursive iteration converges [13].

3.3. Computational Issues

Considering the size of the system in a typical reconstruction problem, the implementation has to meet stringent computational constraints in order to make the method practical. It is important to precompute the matrices A_1, \dots, A_{n-1} using relation (5) and by exploiting the sparsity and symmetry of A_i and the special structure of U_i . In that case, for the 2D multiresolution scheme described in this section, the computational cost of computing the matrices for all resolution will be in $\mathcal{O}(2^n \times 2^n) = \mathcal{O}(N_{\text{Pixel}})$, where N_{Pixel} is the number of grid points. For example, in the case of linear splines, it is possible to compute the matrices with around $20(2^n \times 2^n)$ multiplications. The matrices are actually non-symmetric if symmetric boundary conditions are used. However, it is possible to work with equivalent modified matrices that are symmetric. This modification has to be incorporated in the sparse computation involved in (5). This is important in order to keep the total memory requirement affordable. Since the matrices are very large, efficient representation of the matrices is also important.

4. EXPERIMENTS

Figure 1 gives experimental results for the reconstruction of an image using linear spline. The subsampling was done by thresholding the gradient magnitude retaining 30 percent of the samples. The reconstructed image size is 257×257

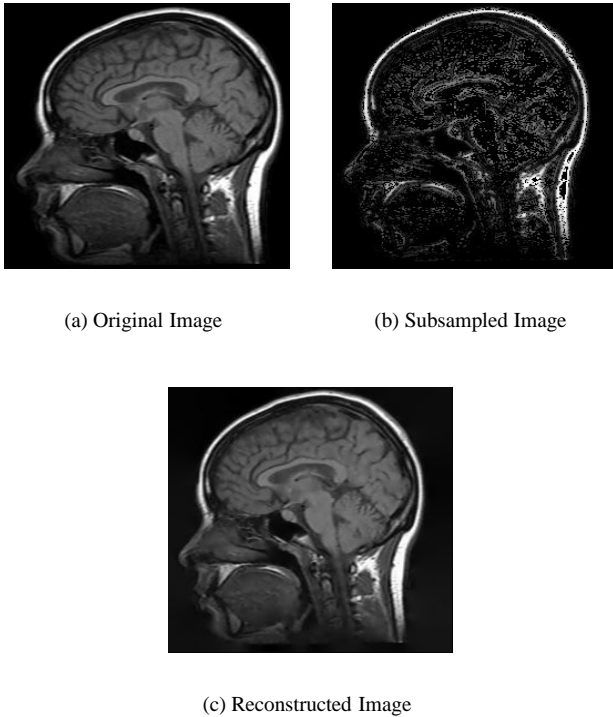


Fig. 1. Image reconstruction from gradient points

and the time taken for reconstruction was 0.5 seconds in Macintosh G4 system, which is quite impressive given the size of the underlying system of equations. Empirically, we find that our algorithm is capable of computing the solution in $\mathcal{O}(N_{\text{Pixel}})$ operations. This is orders of magnitude better than the traditional RBFs methods (thin plate splines). The key point here is that the cost is independent of the number of initial sample points.

5. CONCLUSION

We have developed a variational method for signal reconstruction from irregularly spaced samples on a uniform spline grid. The method allows trade-off between smoothness and reconstruction error (parameter λ), which is particularly useful in the presence of noise. In addition, due its efficient multiresolution and multigrid implementation, it also allows trade-off between computational complexity and reconstruction accuracy (grid size h).

6. REFERENCES

[1] C. D. Barray, C. P. Allot, N. W. John, D. M. Mellor, P. A. Arundel, D. S. Thomson, and J. C. Waterton, "Three dimensional free hand ultrasound: Image

reconstruction and volume analysis," *Ultrasound in medicine and biology*, vol. Vol. 23, No. 8, pp. 1209–1224, 1997.

- [2] J. Duchon, "Splines minimizing rotation-invariant semi-norms in sobolev spaces," *Constructive Theory of Functions of Several Variables*, pp. 85–100, 1977.
- [3] C. A. Micchelli, "Interpolation of scattered data: distance matrices and conditionally positive definite functions," *Constructive Approximation*, vol. Vol. 2, pp. 11–22, 1986.
- [4] J. Meinguet, "Multivariate interpolation at arbitrary points made simple," *Zeitschrift fur Angewandte Mathematik und Physik*, vol. Vol. 30, pp. 292–304, 1979.
- [5] F. L. Bookstein, *Morphometric Tools for Landmark Data*, Cambridge University Press, 1991.
- [6] W. R. Fright J. C. Carr and R. K. Beatson, "Surface interpolation with radial basis function for medical imaging," *IEEE Transactions on Medical Imaging*, vol. Vol. 16, No. 1, pp. 96–107, 1997.
- [7] W. A. Light R. K. Beatson and S. Billings, "Fast solutions of radial basis function interpolation equations: Domain decomposition methods," *SIAM Journal of Scientific Computing*, vol. Vol. 22, No. 5, pp. 1717–1740, 2000.
- [8] C. de Boor, *A Practical Guide to Splines*, Springer-Verlag, 1978.
- [9] Michael Unser, "Splines: A perfect fit for signal and image processing," *IEEE Signal Processing Magazine*, vol. Vol. 16, No. 2, pp. 22–38, Nov. 1999.
- [10] Michael Unser, "Multigrid adaptive image processing," *IEEE Internation Conference on Image Processing*, vol. Vol. 1, pp. 49–52, Oct. 1995.
- [11] M.-H. Yaou and W.-T. Chang, "Wavlet transform in scattering data interpolation," *Electronic Letters*, vol. Vol. 29, No. 21, pp. 1835–1837, Oct. 1993.
- [12] Richard Szeliski, "Fast surface interpolation using hierarchical basis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. Vol. 12, No. 6, pp. 513–528, June 1990.
- [13] W. Hakbusch, *Iterative Solutions of Large Sparse Systems of Equations*, Springer-Verlag, 1994.
- [14] V. E. Henson W. L. Briggs and S. F. McCormick, *A Multigrid Tutorial*, SIAM, 2000.