Global-Scale Three-Dimensional Statistical Estimation

Paul W. Fieguth *

Department of Systems Design Engineering
University of Waterloo
Waterloo, Ontario, Canada
pfieguth@uwaterloo.ca

Abstract

In the statistical estimation context, the difficulty in the transition from two– to three-dimensional problems is *much* more than an increase in the number of pixels: the complexity *per pixel* is cubed!, even for efficient estimation algorithms. That is, even aside from the other challenges posed by 3D estimation (e.g., complex data structures, the difficulty of determining empirical statistics), the computational issues *alone* are significant and merit attention. This paper introduces an alternative approach, motivated by multiscale estimation and the multipole algorithm of mathematical physics.

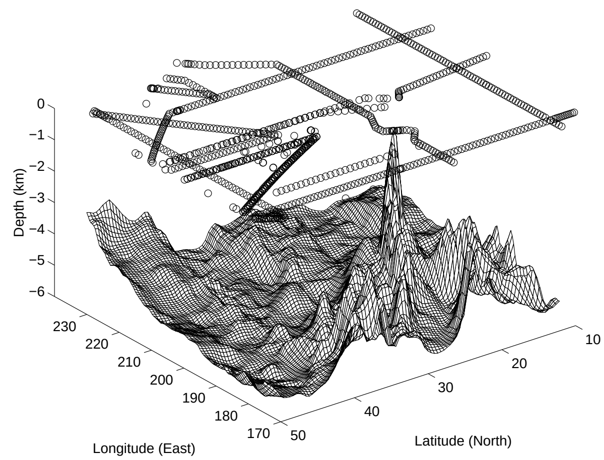

Figure 1: The Pacific basin bathymetry and sparse measurement distribution (measurements are taken at all depths below each circle).

1 Introduction

The statistical estimation of even modestly-sized three-dimensional and large, global scale, two-dimensional remote sensing problems problems presents tremendous and pertinent challenges: heightened environmental awareness and concerns have led to an explosion in the quantity of remotely-sensed data, much of which contains irregular gaps and nonstationary underlying fields. Figure 1 shows one motivating example: the estimation of the volumetric temperature distribution in the central Pacific based on sparse data[2].

The origin of the difficulty in producing statistical estimates is straightforward. Methods such as nested dissection or multiscale estimation[1] are all based on a recursive divide-and-conquer metaphor: a subset of the random field needs to be found, such that conditioned on this subset the remaining portions of the field are conditionally independent and hence can be processed separately. For example, the four quadrants of a first-order Markov random field can be decorrelated by conditioning on the boundary pixels, shown in Figure 2. So whereas a *single* pixel can decorrelate the two halves of a one-dimensional process, a *column* of pixels is required for the 2D field in Figure 2, and a whole *plane* of pixels in three dimensions (Figure 1). Thus for an $n \times n \times \ldots$ hypercube of voxels in d dimensions, the computational effort to solve the estimation problem is

$$\begin{aligned} & \mathcal{O}(n^{(d-1)^3}) && (1) \\ = \quad & \mathcal{O}(n^{(2d-3)}) \text{ per pixel} && (2) \end{aligned}$$

in other words, $\mathcal{O}(n^3)$ *per pixel* in 3D, which very rapidly becomes infeasible for large n. That is, aside from the other challenges posed by estimation in three-dimensions, the computational issues *alone* are significant and serve as the focus of the research in this paper.

*The research of this paper was supported in part by the Natural Science & Engineering Research Council of Canada, and by the Office of Naval Research under Grant N0014-91-J-1004.

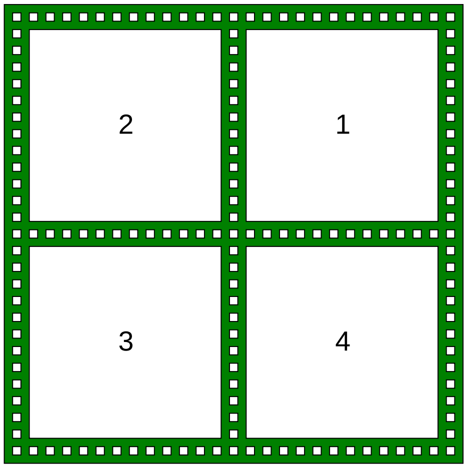

Figure 2: Densely sampled boundaries which conditionally decorrelate the four quadrants of a first-order Markov random field.

2 Multiscale Estimation

The multiscale statistical estimation method[1, 3] has, at its core, the following statistical model:

$$x(s) = A(s)x(s\bar{\gamma}) + B(s)w(s) \tag{3}$$

where s is an index on a tree with parent $s\bar{\gamma}$, A and B are deterministic matrices, and w is a white-noise process. This equation is essentially a restatement of the conditional decorrelation discussed in the Introduction: the whiteness of w implies that the state $x(s\bar{\gamma})$ must conditionally decorrelate all states connected to $s\bar{\gamma}$, leading to a large state dimension or, in the case of reduced-order states (which do not reduce the asymptotic computational complexity), some level of inconsistency across the quadrant boundaries.

For example, Figure 3 shows a measured "tree grain" Markov random field texture. A good reduced-order multiscale state definition[3] was used to produce the estimates in Figure 4, which faithfully (in an MSE sense) capture the random field's statistics, but also manifest artifacts at the quadrant-boundary. A similar approach[2] can be used to solve the 3D problem of Figure 1 by applying the singular value decomposition to the vertical statistics, approximately decoupling the 3D problem into a set of 2D ones, each of which can be solved independently. Although this approach is feasible — sample estimates are shown in Figure 5 — the decoupling makes certain stationarity assumptions and offers limited vertical resolution.

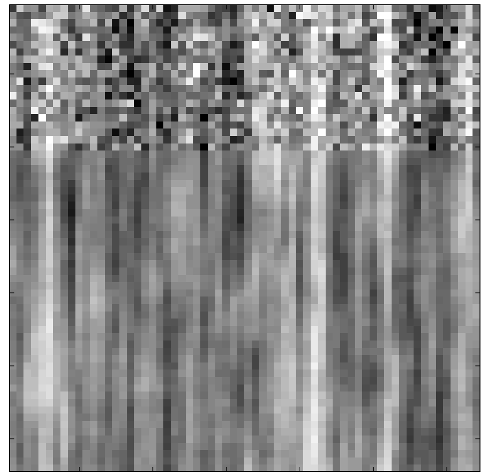

Figure 3: "Tree grain" sample path; the whole field is observed at 0dB SNR, a portion of which is shown at the top of the image.

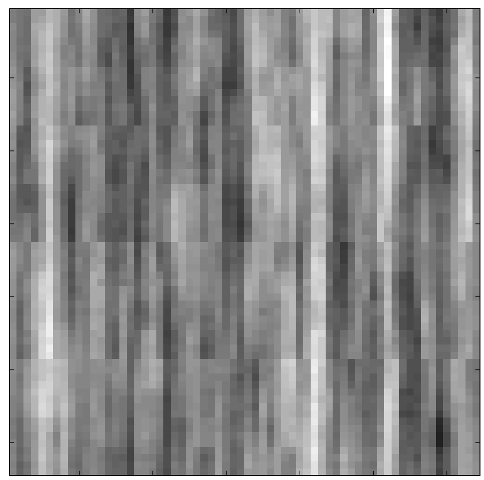

Figure 4: Estimates based on a reduced-order multiscale model applied to Figure 3.

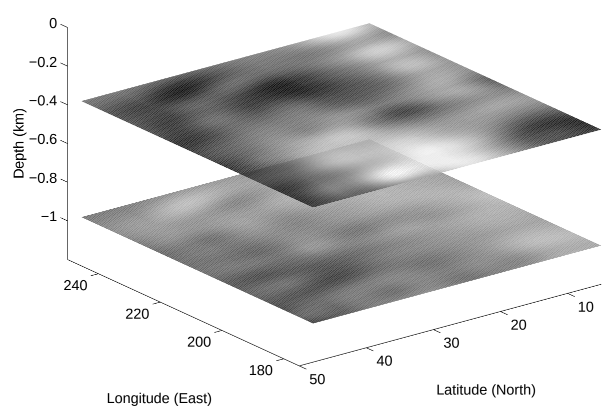

Figure 5: Two planes of Pacific temperature estimates, corresponding to Figure 1, based on a decoupled set of models.

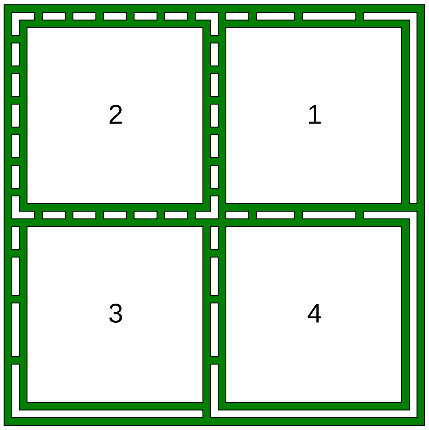

Figure 6: A reduction in state from Figure 2, appropriate for estimating the top-left quadrant.

3 Proposed Approach

The key problem with the state representation in Figure 2 is that it attempts to decorrelate too much: in terms of estimating quadrant "2", keeping details of the distant part of quadrant "4" is largely irrelevant; that is, the reduced state of Figure 6 will perform very nearly as well. Although we now require four such reduced models (one for each quadrant), the computational effort is less than one fourth of that before, giving an overall performance increase. We can apply this insight recursively, for example finding an even further reduced set of states to estimate just a portion of quadrant '2' etc. Taken to the extreme, we can imagine defining a state to estimate just a single pixel; Figure 7 shows an example of such a state, in which the resolution at which each portion of the field is represented is proportional to the distance to the pixel being estimated. Although such an approach may seem impractical, the asymptotic state dimension is only $\mathcal{O}(\log n)$ in 2D *and* in 3D, implying a matrix-inversion effort of only $\mathcal{O}[(\log n)^3]$ per pixel, which is an enormous improvement over (2).

This approach is illustrated in Figure 8, where the "tree grain" field is estimated pixel-by-pixel based on the measurements of Figure 3. The estimates are reasonable, possess no quadrant-boundary artifacts, and the corresponding error statistics (not shown) are accurate to within 10%.

Significant challenges remain. The per-pixel complexity of $\mathcal{O}[(\log n)^3]$ represents matrix-inversion time only; the effort required to compute all of the needed joint statistics can very easily overwhelm this complexity unless we are very careful. Furthermore to apply the concepts of Figure 7 to large 3D problems will require an efficient software implementation. Both issues will be discussed in detail at IMDSP'98.

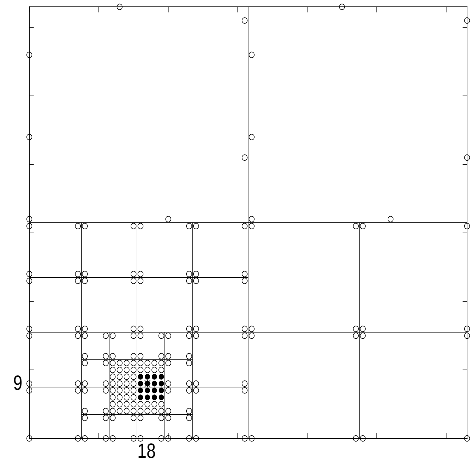

Figure 7: One possible limiting case of Figure 6 for estimating the single pixel (18,9). The state elements are either individual pixels (solid circles) or local averages of pixels (open circles).

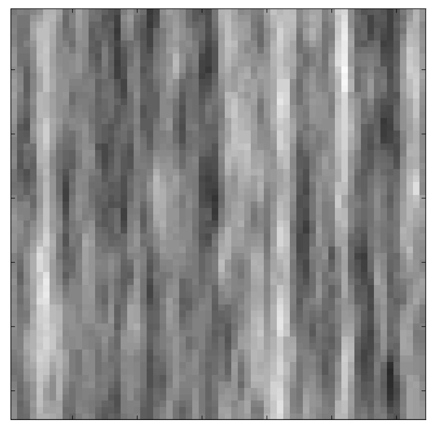

Figure 8: Estimates of Figure 3, but now based on a pixel-by-pixel state assignment.

References

[1] P. Fieguth, W. Karl, A. Willsky, C. Wunsch, "Multiresolution Optimal Interpolation ...", *IEEE GRS* (33) #2, pp.280–292, 1995

[2] D. Menemenlis, P. Fieguth, C. Wunsch, A. Willsky, "Adaptation of a Fast Optimal Interpolation Algorithm ...", *JGR* (102) #C5, pp.10573–10584, 1997

[3] W. Irving, P. Fieguth, A. Willsky, "An Overlapping Tree Approach ...", *IEEE Trans. on Image Proc.* (6) #11, pp.1517–1529, 1997