Efficient Interpolation of Large Image Sequences

F.M. Khellah P.W. Fieguth
University of Waterloo
Waterloo, Ont., N2L-3G1, Canada

Abstract – Dynamic estimation of large-scale remote-sensing image sequences is important in a variety of scientific applications. However, the growing size of such sensed images makes conventional dynamic estimation methods, for example the Kalman and related filters, impractical. In this paper we present an approach that emulates the Kalman filter, but with considerably reduced computational and storage requirements. Our approach is illustrated in the context of a large (512×512) **image sequence of ocean surface temperature.**

I. Introduction

There is a tremendous interest in the processing of image sequences! Although a great deal of this interest stems from the excitement surrounding up-and-coming areas such as *multimedia* and *Internet video*, much of the research is being driven by the fact that storage, bandwidth, and computational power have, for the first time, really increased to the point where it is practical to process sequences of huge arrays of numbers, or images.

Solutions to the image sequence processing problem generally fall into the generic prediction-update structure (i.e., Kalman filter)[3], [10], [13], the general form of which is sketched in Figure 1: a sequence of observed images $\mathbf{y}(\mathbf{t})$ is processed, predicted estimates $\hat{\mathbf{x}}(t|t-1)$ are inferred from an estimated motion field $\mathbf{m_t}$, and the updated estimates $\hat{\mathbf{x}}(t|t)$ are driven by a residual field $\nu(\mathbf{t})$ which encodes the information present in $\mathbf{y}(\mathbf{t})$ which is not contained in $\hat{\mathbf{x}}(t|t-1)$. However the Kalman filter is totally impractical for large remotely-sensed images.

Our research stems from the desire to estimate sea-surface temperature (SST) from the sparse measurements taken by the Along Track Scanning Radiometer (ATSR)[12]. The goal is to produce dense images of SST with associated estimation error variances, but preserving the local features which are present in the measurements. Because of the extreme sparsity in the data, and because ocean SST is highly correlated from one day to the next, clearly the only practical approach is to perform smoothing over time; that is, to process the time-dependent sequence of measurement images. The specific challenges in dynamic estimation are (a) the propagation/storage of huge error covariances in prediction and (b) the inversion of large matrices in update.

II. Prediction

Suppose we have some discretized dynamic model A, such that

$$\underline{x}(t+1) = A\underline{x}(t) + \underline{w}(t). \qquad (1)$$

In general, the exact state prediction according to (1) is straightforward. Furthermore, for sparse, stationary dynamics A can be represented implicitly by a kernel, allowing the exact state prediction to be computed as a convolution

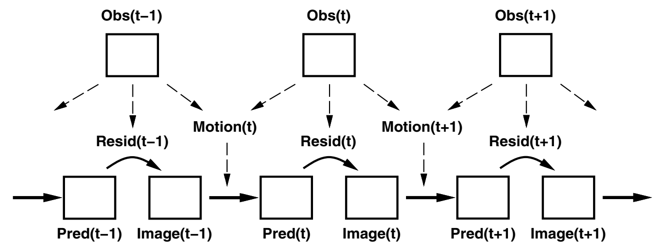

Fig. 1. Standard prediction-update process: An image sequence can be used to infer motion information and a residual, which determine the time-to-time prediction and update. A wide variety of problems can be cast into this framework.

$$\hat{\underline{x}}(t+1|t) = A_k * \hat{\underline{x}}(t|t). \qquad (2)$$

The much greater challenge lies in the prediction of the error statistics, in which the estimation error is propagated through time as

$$\tilde{P}(t+1|t) = A\tilde{P}(t|t)A^T + Q. \qquad (3)$$

For large-scale dynamic estimation problems the exact, brute-force computation of (3) is impossible. Our approach for the error prediction step is to parameterize the error covariances $\tilde{P}(t|t), \tilde{P}(t+1|t)$. Note that any positive semi-definite matrix P can be written as

$$P \quad = \quad \{\underline{p}\underline{p}^T\}^{\frac{1}{2}} \odot \Phi \qquad (4)$$

where $\odot$ refers to element-by-element multiplication, and where Φ is the matrix of normalized correlations:

$$\begin{bmatrix} 1 & \rho_{(1,1)(2,1)} & \cdots & \rho_{(1,1)(n,1)} \\ \rho_{(2,1)(1,1)} & 1 & & \vdots \\ \vdots & & \ddots & \\ \rho_{(1,1)(n,1)} & \cdots & & 1 \end{bmatrix} \qquad (5)$$

The statistics prediction problem then reduces to the prediction of the parameters $\tilde{\underline{p}}(t+1|t), \Phi(t+1|t)$. To ensure the positive-definiteness of $\Phi(t+1|t)$ we model the correlations as exponential. The relationships between the error variances and correlation lengths is then found empirically, as shown in Figure 2.

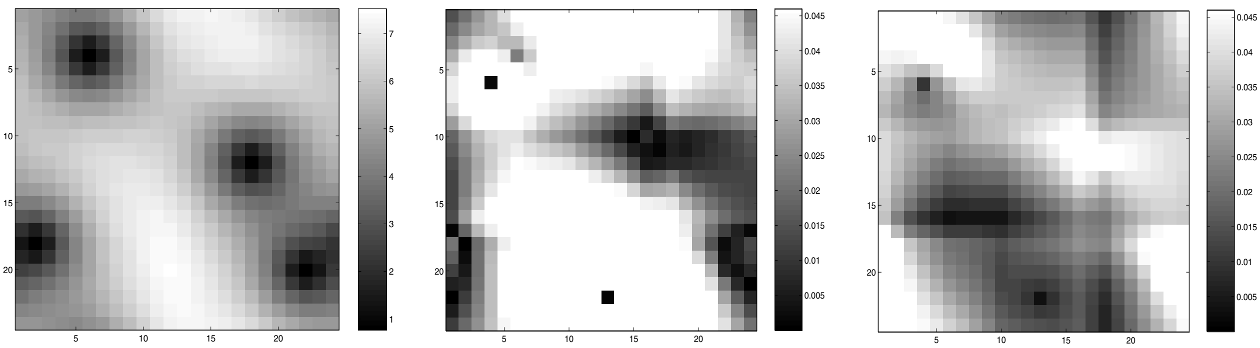

(a) (b) (c)

Fig. 3. Fixed correlation-length prior models: Suppose we wish to fuse two new measurements with a nonstationary prior model (a), where the darkness shows the correlation length. The inconsistency between the exact estimates and estimates are based on a correlation length of (b) $l_1 = 0.6$ and (c) $l_2 = 9$, where dark shades represent higher quality. Note that although both (b),(c) contain regions of significant error, an excellent estimated image could be found through the selection of appropriate subsets of (b) and (c).

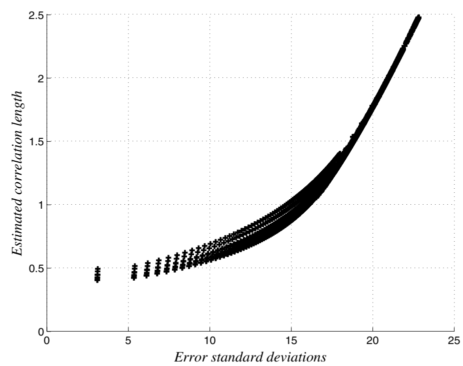

Fig. 2. The empirical relationship between the error standard deviations $\sqrt{p}$ and the correlation length L for diffusion processes with process correlations lengths of (a) 4.0 and (b) 12.0.

III. Update

Many efficient estimators[1], [4], [5], [8], [9] apply most easily to *static* problems. Challenges arise from *dynamic* estimation [9], in which the error statistics become non-stationary. Even our proposed parameterization has problems, in that the exponential correlations are guaranteed to be positive definite only for a fixed correlation length, not a space-varying one. However a fixed correlation length may be wrong and will lead to incorrect estimation results in various regions of the 2D process. For example, suppose we want to update a nonstationary prior model with new measurements; Figure 3 plots the quality of the estimates for two different correlation lengths. Clearly the estimates computed using a short correlation length, Figure 3(b), are more accurate in those parts of the prior having a short correlation length, and similarly Figure 3(c) for longer lengths.

The key insight is that although both Figure 3(b),(c) contain regions of significant error, an excellent estimated image could be found through the selection of appropriate *subsets* of (b) and (c). Even better, can we interpolate the two sets of estimates to acquire even better estimates for intermediate correlation lengths? Finally, a key problem is how to define a nonstationary prior covariance which is guaranteed to be positive definite, a problem not yet fully solved in two dimensions.[11], [9]

TABLE I

Maximum error as a function of K and process correlation length.

Process correlation length	*Number of priors K* 3	5	7	9
3	0.0109	0.0013	0.0004	0.0002
8	0.0384	0.0070	0.0022	0.0009
17	0.0789	0.0145	0.0045	0.0019
25	0.0996	0.0183	0.0056	0.0024

We address all of these issues by finding the estimates and error variances as the nonstationary interpolation

$$\hat{\hat{x}}(i,j) = \sum_{k=1}^{K} \alpha_x(i,j,l_k)\,\hat{x}(i,j,l_k) \qquad (6)$$

$$\hat{\tilde{p}}(i,j) = \sum_{k=1}^{K} \alpha_p(i,j,l_k)\,\tilde{p}(i,j,l_k) \qquad (7)$$

of an ensemble of K stationary estimation problems, each of which is easily guaranteed to be positive definite and can be solved efficiently. The interpolation weights α are found through least-squares, an example of which is shown in Figure 4, and an optimization criterion is defined in order to determine *which* stationary problems to solve. The result is an elegant solution with a very high degree of accuracy.

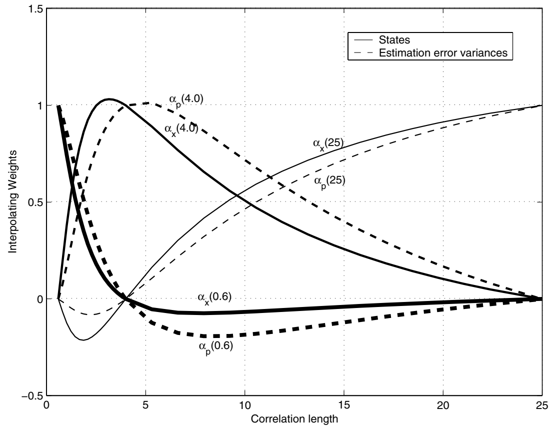

Fig. 4. The general shape of the interpolating weights for the states α_x (solid lines) and the estimation error variances α_p (dashed lines)

IV. Results

Table I shows how the maximum error in the update step varies with the number of stationary models K to interpolate. Even a modest number of models admit extremely accurate updates, and virtually all of the error in the dynamic estimates can be traced to the prediction.

A sample of the dynamic estimation results obtained by the proposed estimator, applied to five months of ATSR data, is shown in Figure 5, showing estimates and corresponding estimation error standard deviations.

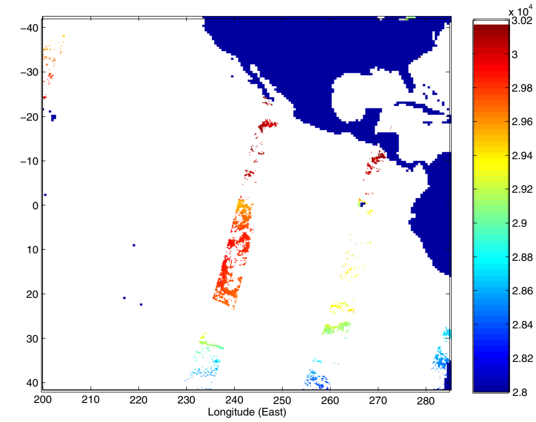

(a) Sample ATSR Measurements

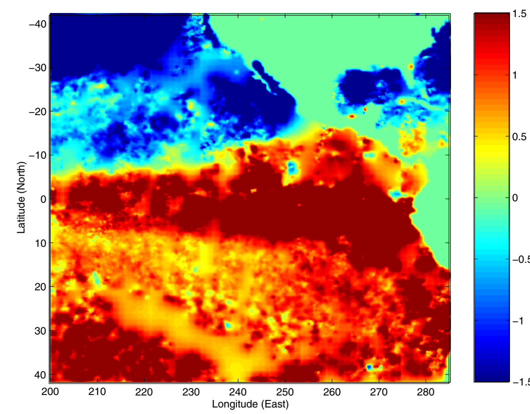

(b) Anomaly Estimates (K)

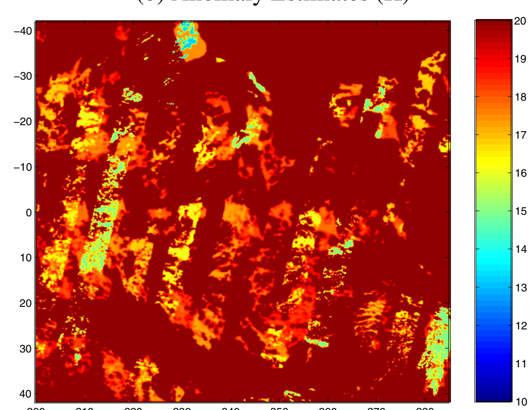

(c) Estimation Error Std (K)

Fig. 5. Reconstruction of dynamic sea surface temperature over the period of June to October, 1992.

References

[1] Amir Asif and Jose Moura. Data assimilation in large time-varying multidimensional fields. *IEEE Transactions on Image Processing*, Vol. 8, No. 11:1593–1607, 1999.

[2] M. A. Cane, A. Kaplan, R. N. Miller, B. Tang, E. C. Hackert, and A. J. Busalacchi. Mapping tropical Pacific sea level: Data assimilation via a reduced state space Kalman filter. *JGR*, 101:22,599–22,617, 1996.

[3] T. Chin. On Kalman filter solution of space-time interpolation. *IEEE Transactions On Image Processing*, Vol. 10, No. 4:663–666, 2001.

[4] T. Chin, W. Karl, and A. Willsky. Sequential filtering for multi-frame visual reconstruction. *Signal Processing*, Vol. 28:311–333, 1992.

[5] Kenneth Chou, Alan Willsky, and Albert Benvensite. Multiscale recursive estimation, data fusion, and regularization. *IEEE Trans on Automatic Control*, Vol. 39, No. 3:464–477, 1994.

[6] Jerome Coleman. *Gaussian Spacetime Models: Markov Field Properties*. PhD thesis, University of California at Davis, David, CA, 1995.

[7] William J. Emery. The advanced very high resolution radiometer (AVHRR): A brief reference guide. *Photogrammetric Engineering and Remote Sensing*, Vol. 58, No. 8:1183–1188, 1992.

[8] Paul Fieguth, William Karl, and Alan Willsky. Multiresolution optimal interpolation and statistical analysis of topex/poseidon satellite altimetry. *IEEE Trans. on Geoscience and Remote Sensing*, Vol. 33, No. 2:280–292, 1995.

[9] Terrence Ho, Paul Fieguth, and Alan Willsky. Computationally efficient steady-state multiscale estimation for 1-d diffusion processes. *Automatica*, Vol. 37:325–340, 2001.

[10] J. Kim and J. Woods. Spatio-temporal adapive 3-D Kalman filter for video. *IEEE Trans. on Image Process*, Vol. 6, No. 3:414–424, 1997.

[11] H. Lev-Ari, S. Parker, and T. Kailath. Multidimensional. maximum-entropy covariance extension. *IEEE Trans. Information Theory*, Vol. 35:497–508, 1989.

[12] C. Mutlow and A. Zavody. Sea surface temperature measurements by the along -track scanning radiometer on the ERS1 satellite: Early results. *Journal of Geophysical Research*, Vol. 99, No. C11:22575–22588, 1994.

[13] J. Woods and C. Radewan. Kalman filtering in two dimensions. *IEEE Trans. Inform. Theory*, Vol. IT-23:437–481, 1977.