MIT Libraries | DSpace@MIT

MIT Open Access Articles

Quantifying Uncertainty for Coherent Structures

Massachusetts Institute of Technology

Quantifying Uncertainty for Coherent Structures

S. Ravela*

Earth Signals and Systems Group
Earth, Atmospheric and Planetary Sciences
Massachusetts Institute of Technology
54-1818, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

Abstract

Field Alignment is a useful and often necessary preprocessing step in contemporary geophysical state and parameter estimation of coherent structures. In an advance, we introduce a new framework for using Field Alignment to quantify uncertainty from an ensemble of coherent structures. Our method, called *Coalescence*, discovers the mean field under non-trivial misalignments of fields with complex shapes, which is especially difficult to calculate in the presence of sparse observations. We solve the associated Field Alignment problem using novel constraints derived from turbulent displacement spectra. In conjunction with a continuation method called Scale Cascaded Alignment (SCA), we are able to extract simpler explanations of the error between fields before cascading to more complex deformation solutions. For coherent structures, SCA and Coalescence have the potential to change the way uncertainty is quantified and data is assimilated. We illustrate utility here in a Nowcasting application.

Keywords:
Data Assimilation, Coherent Structures, Field Alignment, Scale-Cascaded Alignment, Coalescence, Nowcasting, Storm Prediction.

1. Introduction

Environmental data assimilation is the methodology for combining imperfect model predictions with uncertain data in a way that acknowledges their respective uncertainties. It plays a fundamental role in DDDAS as the mechanism by which environmental models are updated. However, data assimilation can only work when the estimation process properly represents all sources of error. The difficulties created by improperly represented errors are particularly apparent in mesoscale meteorological phenomena such as thunderstorms, squall-lines, wild fires, hurricanes, precipitation, and fronts. They also show up elsewhere, for example in reservoir applications where permeability fields have strong contrast.

Errors in mesoscale models can arise in many ways but they often manifest themselves as errors in position. We typically cannot attribute position error to a single source and it is likely that they are the aggregate result of errors

---

Telephone: +16172530997

Email address: ravela@mit.edu (S. Ravela)

in parameter values, initial conditions, boundary conditions and others. In the context of cyclones, operational fore-casters resort to ad hoc procedures such as bogussing [1]. A more sophisticated alternative is to use data assimilation methods. Unfortunately, sequential [2], ensemble-based [3] and variational [4] state estimation methods used in data assimilation applications adjust amplitudes to deal with position error. Adjusting amplitudes does not really fix position error, and can produce unacceptably distorted estimates instead. As seen in Figure 1, showing a simulated "one dimensional front," the true signal (dashed line) observed sparsely (blue dots) by a noisy sensor produces a reasonable estimate (red) from an initial guess (green) when the position error is negligible. It causes disaster (right image) when the error is relatively larger.
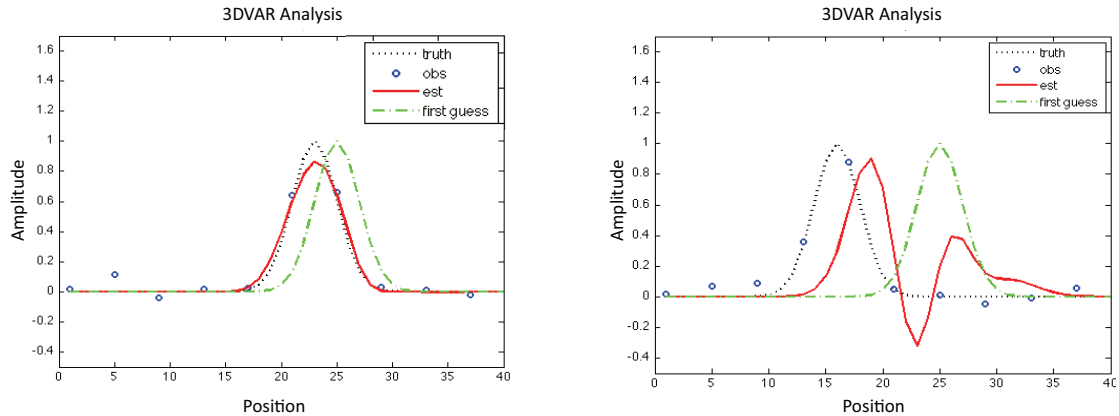

Figure 1: Estimation with "3DVAR" or ensemble methods works well when position errors are small (left). The analysis can be seriously distorted when these errors are large (right).

Indeed, a growing number of researchers are pointing to the role that position errors play in spatial estimation [5, 6, 7]. The issue captured by Figure 1 appears to occur irrespective of the nature of the contemporary estimation procedure used. Solutions have also been proposed, most notably as an optimization problem expressed in the space of grid positions and field amplitudes, with suitable constraints on the grid deformation (or pixel displacement) [8, 5, 9]. Rather than just adjusting amplitudes, this is a more general way of minimizing misfit between two spatial fields. Not surprisingly, it can be applied to many geophysical, meteorological and oceanographic problems [8, 10, 11, 5, 12, 13, 14, 15, 9]. In Section 2, we will summarize the advances.


Figure 2: Smoothness constraint in alignment formulations is common practice. In the top row, a template (left image) is being aligned to the target (right image), which is a rotated and translated version. The iterations of the optimization produce a complex deformation sequence. Below, SCA, a new proposed method, automatically recovers the dominant modes of deformation by seeking the simplest explanation of the deformation first.

However, there are many outstanding issues. This paper will focus on two, progressing from previous DDDAS

work [8]. The first issue concerns the appropriateness of constraints on grid deformation. Most methods rely on smoothness constraints or parameterizations such as splines. Such methods are limited because even as they display remarkable capabilities at "morphing" one pattern into another, such *invariance* comes at the cost of *selectivity*. That is, the more complex a dissimilarity in shape between fields that an algorithm can undo, the more difficult it is to control the deformation solutions produced. In Figure 2, for example, we see that a "cross" is merely translated and rotated (see first and fifth columns). However, snapshots of the deformation solution with smoothness constraints (top row) depicts a complex "viscous" motion. This is not surprising but it is not appropriate either. How might we trust such a method when observations are sparse? We propose a new method, Scale Cascaded Alignment (SCA) [9, 11, 10], that returns the simplest possible "explanation" (see bottom row). The improved control of the deformation solution (selectivity) without loss of a "morphing" ability (invariance) will, as we shall see, turn out to be a key benefit. SCA is discussed in Section 3.


Figure 3: An ensemble of one-dimensional fronts on the left with position and amplitude errors. The ensemble mean is an inappropriate mean field, producing a smeared bimodal front that loses the frontal shape. We introduce a method called Coalescence, which calculates the mean by marginalizing relative position errors in the ensemble (right).

The second issue considered is uncertainty quantification for coherent structures where SCA plays a key part. Consider, as an illustration, an ensemble of "one dimensional" fronts that contain position and amplitude errors, as shown in Figure 3 (left). If we were to ask what the mean front is, clearly the simplest solution is to take the mean of these fields, that is, calculate the mean vector. That would be terribly wrong, of course, because the mean simply does not look like any front in the ensemble. It would be terrible, for example, if the mean of an ensemble of hurricane fields produced a mean field with multiple vortices! The solution to this problem is not trivial, but is a fundamental step when quantifying uncertainty by moments.

We introduce a technique called Coalescence, wherein the mean amplitude field is calculated by marginalizing relative position errors. The method invokes an "N-body" type solution where each member in the ensemble gravitates to the others. In so doing, all of them discover a mean position where the amplitude mean is meaningful. This technique is an extension of a preliminary approach to this problem [12] and we should note its validity in the presence of complex misalignments, structural shape variability and sparsity. Coalescence is discussed in Section 4.

A large number of applications benefit from these new developments including data assimilation [5], verification [7], reservoir modeling, pattern recognition [9], and velocimetry [11, 10]. We will present storm tracking in Section 5.

2. Methodology

Let us consider a gridded scalar field X, and a deformation field (or displacement field) q, which is a vector field, and $X \circ q$ as the deformation of X by q, as shown in Figure 4. We think about a second field Y, which may be gridded

$$\mathbf{X} \qquad \mathbf{q} \qquad \mathbf{X(r - q(r))}$$

Figure 4: A gridded field X is deformed by applying a displacement vector field q to its coordinates r

or not, and related to the first field by the observation operator h. From a suitable Bayesian formulation, one may synthesize the following objective [5, 9]:

$$J(X, q) = \left[X - X^f \circ q \right]^T B(X^f \circ q)^{-1} \left[X - X^f \circ q \right] + \left[Y - h(X^f \circ q) \right]^T R^{-1} \left[Y - h(X^f \circ q) \right] + \Lambda(q) \qquad (1)$$

The term B in (1) is the background error covariance, q is the displacement field, X^f the forecast, and $h(\cdot)$ the observation operator and Λ the constraints on displacement. Like a pixel-wise amplitude error, the displacement or deformation vector field q is defined at every grid coordinate. Note that (1) can be trivially extended to use an ensemble by writing $J_e(X, q) = \sum_s J(X_s, q_s)$, with $B \doteq B(\{X_s^f \circ q_s\}_{s=1}^S)$ also being derived from the S-member ensemble.

If we assume that the forecast distribution is perfect without position errors and suitably Gaussian, then it is trivial to see (via a Taylor expansion of (1)) that non-zero position errors inflate the forecast error (amplitude) covariance. Although non-Gaussian amplitude errors may result from position errors, we show that the covariance is always weakened (for non-zero displacements) and this is the principal problem plaguing effective estimation [5]. The corrupted amplitude covariance, especially when it is flow-dependent, can no longer achieve the Cramer-Rao bound that was achieved in a world without position errors.

Error between spatial fields are quantified and reduced in position and amplitude subspaces in (1) without explicit detection of feature positions. However, if feature positions are available, they can be easily incorporated as boundary conditions. Because feature positions need not explicitly be specified, situations where sparse observations are present or feature positions are difficult to specify (e.g. fronts, complex storm shapes) can be directly addressed. Because the deformation solution is dense, non-trivial alignments can be easily represented (as opposed to a strictly Lagrangian approach), including local skews, shears, rotations, divergence, convergence and higher modes [5].

The error between spatial fields can be jointly reduced in position and amplitude subspaces optimally *if and only if* the sensitivity of the amplitude error covariance to position errors is accounted, because B is dependent on q [5]. This is a non-trivial problem except when (a) the amplitude error covariance is assumed to be static or (b) when position and amplitude errors are assumed to be independent or (c) jointly Gaussian. These assumptions will produce poor results when reality does not match them (which is almost always), but they are what have been assumed in earlier work [7, 6]. We do not make such an assumption and were the first to show, instead, that the joint error-reduction problem can be solved using sampling (ensemble) methods [5].

We were also the first to show that the error between spatial fields can be reduced in position and amplitude in two steps in a systematic way (though ad-hoc approaches were considered contemporarily), that is they arise by considering the Euler-Lagrange equations of (1) [8, 5]. This approximation yields two numerical processes: a Field Alignment process and an amplitude adjustment process, the latter being familiar as contemporary data assimilation (in 3D). The Field Alignment process becomes a pre-processor, which makes practical application feasible and with broad utility [8, 5].

The Euler-Lagrange equations can be written as follows. First equation, which is iterated in deformation is:

$$
\begin{aligned}
Z_i &\doteq H^T R^{-1}(Y - X_i^f) \\
\tilde{q}_i &\doteq L_i q_i \\
\tilde{q}_i[k] &= \nabla X_i^{fT}[k]\, Z_i[k] = f_i[k]
\end{aligned}
\tag{2}
$$

H is the Jacobian of the observation operator, and q_i is the instantaneous displacement in the i^{th} iteration, $L_i q_i = d\Lambda/dq_i$, $X_0^f = X^f$, $q_0 = 0$ and $X_i^f = X_{i-1}^f \circ q_{i-1}$. Note that (2) is evaluated for each component of $\tilde{q}_i$'s vector space at pixel or grid point k. The second equation, in amplitude, may be written as:

$$
X^a = X_*^f + B_* H^T (H B_* H^T + R)^{-1}(Y - X_*^f)
\tag{3}
$$

$X_*^f \equiv X^f \circ (\sum_i q_i)$ is the final deformed first guess, which may in practice be obtained via the iteration: $X_i^f = X_{i-1}^f \circ q_{i-1}$. $B_* = B(X_*^f)$ is the final covariance, and X^a is a "quadratic update" familiar as OI/3DVAR/KF/EKF/EnKF. Note that (2) and (3) are approximations to a direct solution.

Dynamical Balance is better preserved when data is assimilated by field alignment and the technique defined by (2) is easily extended to multivariate fields and vector fields [8, 5]. However, constraints must be imposed on the nature of the deformation. We showed that constraints on the displacement field can be simple "fluid-like" ones, such as smoothness and non-divergence, see (4), but they need to be specified at an appropriate length scale to be effective, which is often unknown. Otherwise, they produce an inappropriate, aliased explanation of the true deformation. In contrast, we posit that by seeking the simplest explanation [9, 11], simplicity being defined by the smallest spatial wave number (or scale) a deformation is restricted to, useful and robust solutions can be produced. By cascading through deformation scales (from translation to "turbulence"), dominant modes of misalignment can also be estimated. We call this Scale Cascaded Alignment (SCA), discussed next.

3. Scale Cascaded Alignment

With smoothness, and non-divergence, two popular constraints of the Tikhonov type, the iterated alignment equation (2) reduces to:

$$
w_1 \nabla^2 q_i + w_2 \nabla(\nabla \cdot q_i) - f_i = 0
\tag{4}
$$

In what follows, we drop the explicit notation for iteration i and use q to denote the instantaneous displacement. Let us suppose a vector field $q(p) \doteq (q_x(p), q_y(p))^T$ with a Fourier transform pair $Q(\omega) = (Q_x(\omega), Q_y(\omega))^T$. Similarly, let $f \leftrightarrow F$ be a transform pair. We can then write a solution to (4) in wave-number space ($m \neq 0$ and $n \neq 0$) of the form:

$$
\begin{bmatrix} Q_x \\ Q_y \end{bmatrix} =
\begin{bmatrix} H_a & H_b \\ H_b & H_c \end{bmatrix}
\begin{bmatrix} F_x \\ F_y \end{bmatrix}
\tag{5}
$$

For the sake of simplicity, let us set $w_2 = 0$, which leads to Laplace-Beltrami flow (but the problem obviously remains regularized). The resulting filter is the Fourier Transform of the Laplacian (Bi-Laplacian is also commonly used), which prescribes a power-law spectrum for the instantaneous deformation.

$$
H = H_a = H_c = -\frac{1}{w_1 \|\omega\|_2^2}
\tag{6}
$$

The displacement spectrum Q is capable of producing complex deformations because the filter H admits high frequencies. It is capable of aliasing a "simple" cause into a complex explanation, as shown in Figure 2. To control the deformation better, we propose the following three steps:

1. Laplace Approximation of Power Law. In the viscous formulation, w_1 only controls the convergence rate (provided stability is maintained) but not the deformation spectrum. The shape of the filter can be controlled by changing the power law itself, but here we propose a Laplace "distribution" as an approximation: $e(r) = \beta e^{-|r|/2\alpha^2}$, where r is the spectral radius. This formulation overcomes the singularity in the power law and it can reasonably approximate any power law using the parameters β (gain) and α (bandwidth), thus producing instantaneous displacements ranging from translations to turbulence. A comparison of the original filter and the tunable approximation is shown in Figure 5(a) for a power-law of 2 (smoothness).

(a) Filter approximation of power law.

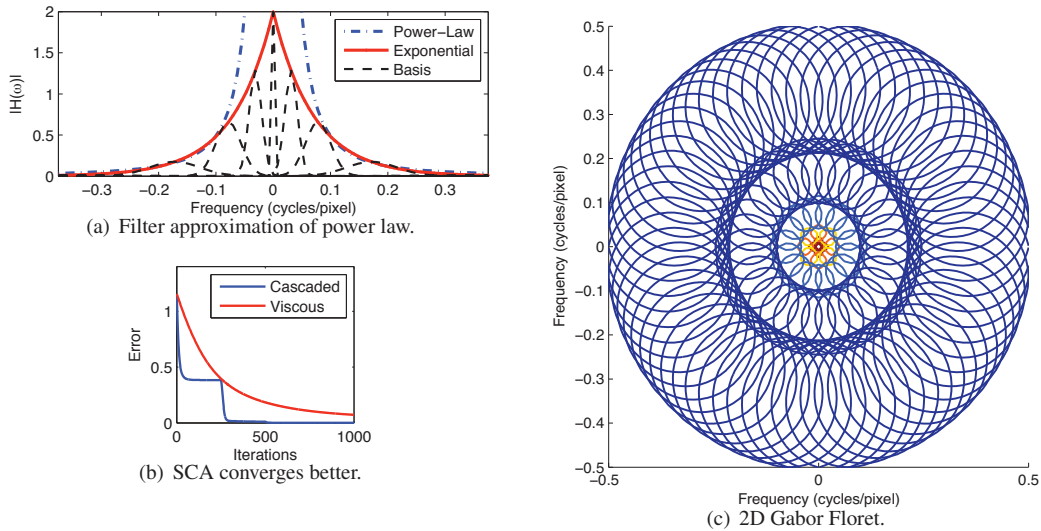(b) SCA converges better.

(c) 2D Gabor Floret.

Figure 5: Smoothness, which yields a power-law, can be approximated by the parameterized Laplace distribution, which may in turn be expressed by a Gabor basis. The resulting algorithm, SCA, performs substantially better.

2. Gabor filter basis for deformations. The deformation spectrum is expressed using a Gabor basis (see Figure 5(b) and Figure 5(c) in 1D and 2D respectively). Thus, we may write the filter as a weighted sum of the basis:

$$H = \sum_i H_i \xi_i \tag{7}$$

The filter H_i is the sum of all filters at a common radius around the spectrum.

3. Scale Cascaded Alignment (SCA). We solve for the total displacement in sequence, starting with filters closest to the origin (DC or translation) and moving higher to more complex wavenumbers. Thus we may iterate with index j as $Q_j = H_i F_j$ before continuing to a finer solution using H_{i+1} and so on, until all the filters organized according to their "spectral radius" are exhausted. In general, filter at wave number 0 is translation, at wavenumber $\frac{1}{2}$ accounts for global deformations, and at the highest wavenumber, total turbulence. Displacement fields corresponding to each filter-bank are returned and their sum is the total displacement that minimizes the model data misfit in position space.

SCA can be implemented as an independent Euler-Lagrange equation, or as a gradient functional in a joint position-amplitude minimization procedure. In the former role, once it converges, the amplitude adjustment described in (3) is performed. In Figure 5(c), we see the reduction in error proposed by SCA versus Smoothness-based Field Alignment, for the example in Figure 2. While "viscous" alignment produces a complex explanation, SCA rapidly converges, returning translation and rotation automatically as the dominant modes. Even as a standalone routine, SCA has numerous applications from weather tracking to pattern recognition and we will review these in Section 5.

4. Coalescence

As noted in Section 1, the calculation of the ensemble mean with position errors is non-trivial. The ensemble of vortices in Figure 6(a), for example, have a mean, shown in Figure 6(b), that looks like no ensemble member at all. It does not account for the fact that the vortices are relatively displaced and not merely realizations of a Gauss-Markov field with a shape-distorted mean field. In Figure 6(c), we *coalesce* under simple deformation constraints. The mean in Figure 6(d) is much more realistic as a consequence. Similarly, in Figure 6(e) and (f), coalesced fronts produce a more representative ensemble mean. We now go on to discuss a general framework for Coalescence.

Figure 6: Coalescence of vortices and fronts produces a better representation of the ensemble mean.

The problem of discovering a mean position is simple if exact correspondence between fields are densely available, but this is impractical in most problems. When correspondences are sparse, or impossible to produce because fields are only sparsely observed, or when the deformations are nonlinear and the fields have complex shapes, the estimation problem becomes much harder.

We devise a procedure that works under these conditions. It estimates the mean field X by quantifying and marginalizing the deformation q_s associated with an ensemble member X_s^f. With a suitable probability density representation, our Bayesian formulation leads to an Expectation-Maximization (EM) problem. In the i^{th} iteration ($i > 0$) of the EM cycle, define $X_{s,i}^f = X_{s,i-1}^f \circ q_{s,i-1}$ as the displaced ensemble member s, and B_i as the ensemble covariance. Omitting, for simplicity, subscript i by writing $X_s^f \doteq X_{s,i}^f$, the E-step becomes:

$$J(X, \{q_s\}|\{X_s^f\}) = \sum_{s=1}^{S} (X - X_s^f \circ q_s)^T B_i^{-1} (X - X_s^f \circ q_s) + \Lambda_s(q_s) \tag{8}$$

The objective J estimates both displacement and the mean (amplitude) field. The displacement constraint $\Lambda_s(q_s)$ can be smoothness, a turbulent spectrum or others. In the M-step, we minimized (8) by treating J as the negative log-likelihood. Looking for a stationary point yields two Euler-Lagrange equations. The first one is simply an expression for the amplitude mean for an S member ensemble:

$$X = \frac{1}{S} \sum_s X_s^f \circ q_s \tag{9}$$

The second equation is one for deformation of each ensemble, and it is solved iteratively in the i^{th} M-step. Allowing the displacement of ensemble member s in iteration j to be $q_s^{(j)}$, the second normal equation is synthesized as:

$$\begin{aligned}
Z_{ts}^{(j)} &\doteq \left[B_i^{-1}(X_t^{f(j)} - X_s^{f(j)}) \right] \\
\tilde{q}_s^{(j)} &\doteq L_s q_s^{(j)} \\
\tilde{q}_s^{(j)}[k] &= \frac{1}{S} \sum_{t \neq s} \nabla X_s^{f(j)T}[k] \, Z_{ts}^{(j)}[k]
\end{aligned} \tag{10}$$

We define $L_s q_s \doteq d\Lambda_s/dq_s$ and the incrementally displaced state vector per iteration as $X_s^{f(j)} \doteq X_s^{f(j-1)} \circ q_s^{(j-1)}$, with $j > 0$ and $q_s^{(0)} = 0$. Please note that B_i is fixed in the i^{th} M-step and (10) is evaluated at each pixel or grid point k for each component of the vector space of $\tilde{q}_s^{(j)}$. Once (10) converges, the solution $q_s^* \doteq \sum_j q_s^{(j)}$ produces a mean field estimate $X_i^* = \frac{1}{S} \sum_s X_s^f \circ q_s^*$. We then return to the next iteration of the E-step by setting $q_{s,i} = q_s^*$ and updating $X_{s,i+1}^f = X_{s,i}^f \circ q_{s,i}$ and calculating B_{i+1}. We reconstruct an objective and thus alternate the E-M steps until convergence. The final estimate of the ensemble is $X_{s,\infty}^f$, the corresponding mean is X_∞ and covariance B_∞.

We may define $B_i = B(\{X_{s,i}^f\})$, that is a simple calculation from the ensemble (and represented in square-root form) at the beginning of the i^{th} step. However, note that $q_{s,0}$ is not defined and B_∞ is the only appropriate covariance, and all prior ones $B_1 \rightarrow B_\infty$ represent a converging sequence. We may, in fact, have a different sequence, one

that accounts for the large uncertainty in knowledge of B_1 and the strong certainty of B_∞ by writing, for example, $B_i = \gamma_i B(\{X_{s,i}^f\}) + (1 - \gamma_i)I$, for a scalar $\gamma_i \in [0, 1]$ and with $i > 0$. We find in current experiments that the EM iterations converge rapidly as most position errors are eliminated in the first M-step iterations. Consequently, we choose $\gamma_1 = 0$ and $\gamma_2 = 1$, but other sequences can also be developed.

An interpretation of (10) is that every ensemble member experiences a "force" from the other ensemble members and in this way, they all gravitate to a "center of mass" or coalesce to a mean (under displacement). To the best of our knowledge, Coalescence has not been introduced this way, but it is a foundation to quantifying ensemble uncertainty beginning with its mean.

5. Nowcasting Storms

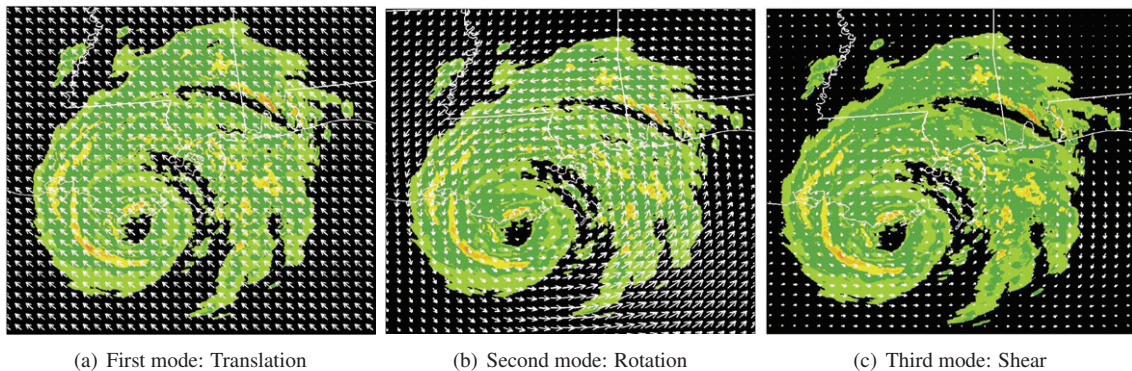| (a) First mode: Translation | (b) Second mode: Rotation | (c) Third mode: Shear |

Figure 7: SCA automatically recovered apparent motion modes for a hurricane from Radar VIL. They correspond strongly to translation, rotation and shear, but without such an explicit parameterization.

Many applications of Coalescence and SCA are emerging. Here, we will discuss their application in a new DDDAS direction: Nowcasting storms.

Weather-related operational delays in US aviation run into the billions of dollars annually. At least 65% of the aircraft delays in the US in 2009 were related to weather. Having effective convective forecasts in the 0-8 hour range can significantly smooth out operational jitters, leading to better managed airspace and traffic [14, 15]. The Nowcasting problem as it is known typically requires forecasting a quantity such as the Vertically Integrated Liquid (VIL) that can be retrieved from radar measurements.

One way to Nowcast VIL is by using a numerical model such as the High Resolution Rapid Refresh model (HRRR). Another way is data-driven. Time series of radar imagery are used to estimate an apparent motion, growth and decay model of significant convective activity (strong VIL) which is used to extrapolate current observations into the future. As it turns out, when "warm-started" with new observations, numerical models display poor skill at short forecast lead times. This is because the adjustments induced by observations have not largely equilibrated with the rest of the model state. Clearly a data-driven VIL forecast will be perfect at a zero lead-time, but its performance decays rapidly in time. The crossover point is typically 3 hours. Therefore, when issuing a 0-8hr forecast, we must blend the data-driven extrapolation model with the numerical model, producing a rendition of truth that is better than either source alone.

Here, we use SCA to build a data-driven apparent motion model, which offers at least one advantage. By extracting the simplest explanation of the apparent motion first and then proceeding to more complex ones at higher wavenumbers, we can capture largescale motion "modes" for producing long term forecasts while rejecting high-frequency motions that tend to be short-lived and introduce undesirable dispersion in the long run. As an example, in Figure 7, we see a snapshot of Hurricane Gustav's Radar VIL, with estimated apparent motion vectors overlaid. These vectors are calculated from radar VIL time series in the immediate past. There is one figure for each mode, increasing in complexity from left to right. The automatically recovered modes by SCA appear to be translation, large-scale rotation and a more localized shear. They correctly summarize the storm's apparent motion to the west-northwest.
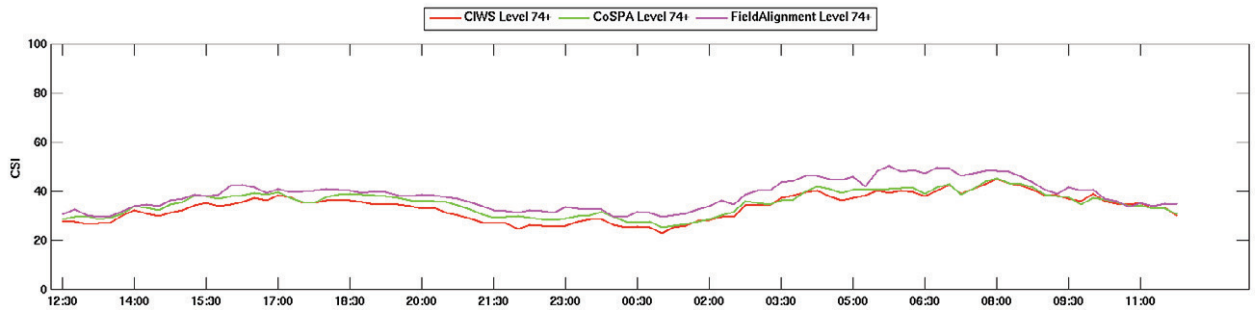
Figure 8: Improvement in Data-Driven two hour forecast skill throughout the day for Hurricane Alex by SCA (pink) over current extrapolation methods.

Thus, we may produce an optimal weighted combination of these modes from a predictability point of view, using it as an "apparent velocity" field with which to advect the current radar VIL into the future. SCA, used this way, outperforms the existing methods in use by Lincoln Laboratory ("LLExtrap") [14, 15], based on correlation Optic Flow, denoted CIWS and CoSPA.

As shown in Figure 8, a 120 minute (two hour) forecast of Hurricane Alex is compared with the verifying radar VIL throughout a diurnal cycle. The CSI[1] scores for SCA is shown in pink, LLExtrap (CIWS[2]) in red and LLExtrap (CoSPA[3]) in blue. The performance is much improved in this case (it is understood by the community that CSI scores generally change very little even when improvements are large). This is promising, and we note that a mechanism to optimize the weighted combination will generalize these results.

The blending problem is posed in general as a fusion of model predictions with extrapolations over a 0-8hr interval, which naturally invokes a two-point boundary value (variational) problem. Coalescence plays a key role within this framework for blending fields at any forecast time. For the experiment depicted here, we use LLExtrap as the data-driven extrapolation, and HRRR as the model. In Figure 9, we see the extrapolation and HRRR forecasts 225 minutes out (middle column). Note that both VIL images have some overlap, but they largely capture different aspects of the storm. The verifying observations are shown in the top-right image. However, since this is a radar image over the ocean, much of the storm is self-occluded, thus producing a small crescent instead of the full storm. The full storm exists, of course, and can be seen from satellite, but depicting brightness temperature (not VIL), shown in bottom-right panel of Figure 9.

We contrast two blending techniques, one produced using the current method and the other by Coalescence (left column). Our reconstructions are more plausible by comparison, whereas the alternative is a little more unrealistic (and flagged by evaluators as a problem case). Coalescence finds a mean field by bringing the two forecasts together in position-space up to local convergence.

Coalescence is turning out to be useful in blending but has additional use in quantifying ensemble variability which will be further elaborated in future work.

Acknowledgements

---

[1]Critical Success Index = hits/(hits+false alarms+misses)

[2]Corridor Integrated Weather System
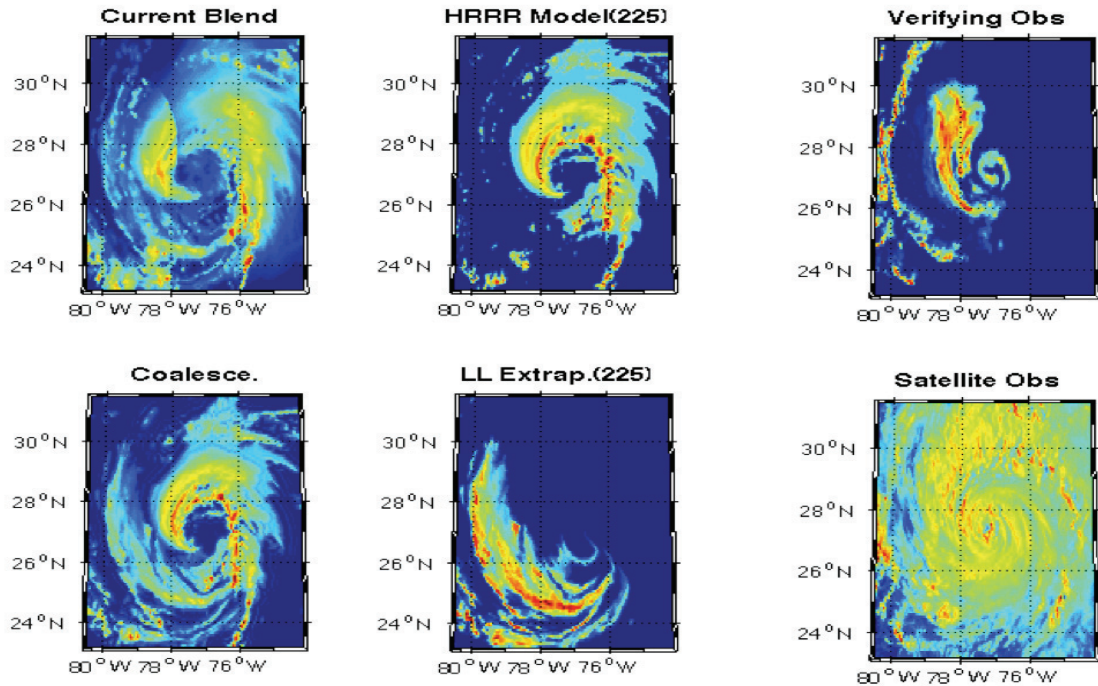
[3]Consolidated Storm Prediction for Aviation

Figure 9: Coalescence is used to blend the outputs of the HRRR model and extrapolation, improving on existing blending methodology.

References

[1] C. Davis, S. Low-Nam, The ncar-afwa tropical cyclone bogussing scheme, Tech. rep., Technical Memorandum, Air Force Weather Agency (AFWA), http://www.mmm.ucar.edu/mm5/mm5v3/tc-report.pdf (2001).

[2] A. Gelb, Applied Optimal Estimation, MIT Press, 1974.

[3] G. Evensen, The ensemble kalman filter: Theoretical formulation and practical implementation, Ocean Dynamics 53 (2003) 342–367.

[4] A. C. Lorenc, Analysis method for numerical weather prediction, Q. J. R. Meteorol. Soc. 112 (1986) 1177–1194.

[5] S. Ravela, K. Emanuel, D. McLaughlin, Data assimilation by field alignment, Physica D 230 (1) (2007) 127–145.

[6] J. D. Beezley, J. Mandel, Morphing ensemble kalman filters, Tellus 60A (2008) 131–140.

[7] R. N. Hoffman, Z. Liu, J. Louis, C. Grassotti, Distortion representation of forecast errors., Monthly Weather Review 123 (1995) 2758–2770.

[8] S. Ravela, Two new directions in data assimilation by field alignment, Lecture Notes in Computer Science, Proc. ICCS 4487 (2007) 1147–1154.

[9] C. Yang, S. Ravela, Spectral control of viscous alignment for deformation invariant image matching, Proceedings of International Conference on Computer Vision 1 (2009) 1303–1310.

[10] S. Ravela, T. Langlois, W. Dupree, M. Wolfson, Data-Driven Weather Tracker Using Scale Cascaded Alignment, US Patent Application 61/442,892 (2011).

[11] S. Ravela, C. Yang, J. William, K. Emanuel, An Objective Framework for Assimilating Coherent Structures, WMO Symposium on Nowcasting (2009).

[12] S. Ravela, V. Chatdarong, How do we deal with position errors in observations and forecasts Annual?, Congress of the European Geosciences Union (2006).

[13] J. K. Williams, Wrf-var implementation for data assimilation experimentation at mit, Master's thesis, Massachusetts Institute of Technology (2008).

[14] M. M. Wolfson, D. A. Clark, Advanced aviation weather forecasts, MIT Lincoln Laboratory Journal 16 (1) (2006) 31–58.

[15] W. J. Dupree, M. M. Wolfson, J. R. J. Johnson, R. A. Boldi, E. B. Mann, K. T. Calden, C. A. Wilson, P. E. Bieringer, B. D. Martin, H. Iskenderian, FAA Tactical Weather Forecasting in the United States National Airspace, Proceedings from the World Weather Research Symposium on Nowcasting and Very Short Term Forecasts (2005).