The Field Alignment System and Testbed

Sai Ravela
Massachusetts Institute of Technology*

August 23, 2013

1 Introduction

The representation of the appearance of coherent structures in fluids is increasingly being recognized as an effective tool for many Geophysical inference problems including Tracking, Nowcasting, Initialization, Verification, Reservoir Modeling, Data Assimilation and Change Analysis, among others. Attention is being drawn to the benefits that position and shape adjustments of coherent structures offer, in addition to grid-point ”amplitude” adjustments, for reducing error between fields and improving parameter estimates. One class of adjustments entail alignment of one field with another, interpreted equivalently as grid deformations or the ’displacement’ of structures apparent in the field. Deformations that reduce error between fields to a convergence (even local) can be invaluable tools to factor out timing and motion errors from the total error and thus highlight growth and decay better and improve verification. Deformations can be useful themselves, most notably as apparent motion models in Tracking and Nowcasting applications. They may be paramount in data assimilation and uncertainty quantification for coherent structures [2].

The notion of deforming or displacing fields for improving analysis is not entirely new. Using correspondence between characteristic visual features to warp and register one field with another, for example, is a well-known verification application. Correlation and variational optic-flow based

*77 Massachusetts Avenue, 54-1818, Cambridge, MA, 02139. Visit http://tinyurl.com/fastmit

approaches for Nowcasting are also well-known. Our work in 2003 initiated use in a Data Assimilation context [4] by noting the obviously poor analyses that result when ”position” information is not considered and one outcome in producing improved analysis was a numerical process called Field Alignment. Our approach offers advantages when one field is only sparsely observed and when complex deformations need to be represented, where parsimonious estimation of deformation is desired, or when the fields are large or multivariate, or when the alignment problem is one of coalescence, i.e. where the reference field must be discovered. Field Alignment allows for ”fluid-like” differential constraints in the search for a deformation minimizing the misfit between fields, or it allows spectral ones that interpret apparent motion in terms of ”turbulent” motion modes. Field Alignment comes with local gradient-based solvers or stochastic versions that produce non-local solutions.

The Field Alignment System and Testbed (FAST) is an ongoing effort at broadening the community around this research. As such, the implementation of various methods along with examples contained in it are hereby made available through the DTC program. In this document, we briefly describe the Field Alignment framework, followed by a description of codes with examples in FAST v3.0, and mechanisms for ongoing updates and upgrades.

2 Framework

Ravela et al.'s field estimation formulation represents a gridded field X[1] deformed by a dense vector field q as $X \circ q = X(p - q(p))$. Using an observation equation with operator h and additive noise n to relate a second field $Y = h(X \circ q) + n$, Ravela et al. consider a Bayesian formulation for estimating the amplitude and deformation as $P(X, q|Y = y) \propto P(Y = y|X, q)P(X|q)P(q)$. In principle, this problem can be solved by Markov Chain Monte Carlo (MCMC) approaches, but there are considerable practical difficulties associated with sampling high-dimensional fluid model

[1]The methodology easily extends to vector fields, scalar field is used here for simplicity.

states (typically $10^5 - 10^6$). By assuming that the distributions are Gaussian but only when the field X^2 is deformation-corrected, and further by assuming that a deformation prior probability distribution can be constructed from displacement constraints, they synthesize the objective [5]:

$$\begin{aligned} J(X,q) &= \frac{1}{2}\left[X \circ q - x^f \circ q\right]^T P(x^f \circ q)^{-1}\left[X \circ q - x^f \circ q\right] \\ &+\frac{1}{2}\left[y - h(X \circ q)\right]^T R^{-1}\left[y - h(X \circ q)\right] \\ &+\Lambda(q) \end{aligned} \quad (1)$$

The matrix P is the amplitude error covariance, which is dependent in general on the deformation q, x^f is the first guess, Λ a potential function on displacement constraints, and R is the observation error covariance. Unreasonable assumptions, namely the independence of position and amplitude errors and a static P have commonly been made to make Equation 1 tractable, which, to be sure is difficult to directly solve due to the dependence of P on q. Ravela et al. show that a solution procedure free from such restrictive assumptions is to use an ensemble representation of the uncertainties [5] whence an solution in terms of iterative least squares [5] or EM becomes available [3]. They also propose another approximate solution that is useful both for deterministic as well as ensemble settings [5, 1, 7]. In this approach, the Euler-Lagrange equations of Equation 1 are used to recover a sequential solution, first for a deformation q, i.e. Field Alignment, and then for an amplitude adjustment that can be implemented using conventional approaches [6].

Only the Field Alignment aspect of Ravela et al.'s formulation is of interest here and with possibly sparse measurements in H. This limited objective can also be obtained from a simplified Bayesian formulation $P(q|X = x^f, Y = y)$ providing a reduced objective

$$J(q|X = x^f) = \frac{1}{2}\left[y - H\ x^f \circ q)\right]^T R^{-1}\left[y - H\ x^f \circ q)\right] + \Lambda(q) \quad (2)$$

[2] Fields and vectors are used interchangeably and conversion between them is implicit in the usage that follows.

Equation 2 when linearized in q lead to the classical optic-flow equations, but Ravela et al., do not perform such linearization. Instead, they formulate the deformation constraint as a scale-parametrized form $\Lambda(q) \equiv \Lambda(q, \sigma)$ that impose a length scale on the smoothness or other differential constraints that the user can tune for a given application. Ravela et al., provide a coarse-to-fine pyramid approach for finding the deformation [1] and they also develop another scale-cascaded representation of displacement constraints that allow for the spectrally most dominant modes of motion to be recovered iteratively [7, 2]. The solution procedure under the differential, pyramid or scale-cascaded versions can be implemented using gradient-based approaches (e.g. BFGS) for the search for a local minima or through a stochastic gradient or MCMC approach to recover deep minima [3].

3 Structure of Testbed Codes

FAST release 3.0 made available to DTC/NOAA consists of several programs for 1D and 2D Field Alignment. The general structure of the first release is as follows:

FA1D/			1D Version
	core/		Core Directory
		codegen/	Library codes accessible through wrapper
		fa1df2cwrap.c	Wrapper
	dat/		Data directory
	examples/		Examples
		example1.m	Matlab version, Stochastic Optimization
		example1.c	C-version
		example1.f90	Fortran version
	startup.m		Matlab startup

FA2D/			2D Version
	core/		Core Directory
		codegen/	Library codes accessible through wrapper
		FA2DNoHF2CWrap.c	Wrapper
	dat/		Data directory
	examples/		Examples
		example0.m	sparse measurement
		example1.m	Radar reflectivity
		example2.m	GFS Example
		example2.c	C-version
		example2.f90	Fortran version
		example3.m	Multivariate example with sparse measurements
		example4.m	Ensemble example
		example5.m	Matlab version
		example6.m	Matlab version
	startup.m		Matlab startup

These codes were built, tested and deployed through out the DTC visit and although this release is not exhaustive list of Field Alignment examples, it contains a core set that enables users to build additional applications. The core routines and examples can be built on most Unix-like platforms and they have been also been built and tested on the Mac. Because the core library is a direct regeneration of Matlab code, testing in Matlab is an easy first step which is followed by the user's application development that can be directly linked to the core library in a ”production” mode. Libraries have been built on local platforms (e.g. Zeus) and a wrapper program (see the directory tables listed here) allows applications to directly run C/Fortran programs and one can study the examples for similar use.

4 Examples

The examples provided cover a range of possible Field Alignment applications. A simple illustrative case is *example0.m*

```
%% A simple 2D example with sparse measurements.
close all;
clear all;
```

```
load example0;  %Some parameters are stored in this file
itermax = 2^10;  %total number of iterations --
mode = 4;  % Bi-Laplacian = 4, Laplacian = 2.
%  4 for stiffer response, 2 when densely observed.
wt = 1/250;  % Scaling. experiment with this value. This value is atypical!
lengthscale = 4;  % sale to 1/4 of cutoff wave number for bilaplacian
graphics = 1;  %return on graphics
debug = 1;  % turn on debugging
%% Call main routine
[qtx,qty,qq,erro] = FA2DC4V2(p1,p2,H,[],[],[],1,1,...
                            itermax,graphics,wt,lengthscale,mode,...
                            debug);
```

The above code simply loads up a case, sets the number of iterations and modes of differential constraints, followed by the weight of the constraint in the objective, and the length scale, which is set at a $1/4$ of the bandwidth of the motion spectrum that an unscaled response would contain. Other experiments will require modification to these parameters and substantial experimentation may sometimes be needed. Please have a look at other examples for use cases.

Another instructive example is *example1.m*, a part of which is shown here:

```
nx = 128; % will be square
    itermax = 2^9; % This should be fixed to 1/2 domain dimensions.
    mode = 2;
    w1 = 1/2; %Make sure solution is stable!
    lscale = 1;
    % No multiresolution alignment in this one--
```

```
[qsavex,qsavey]=FA2DImNoH(m,n,nx,w1,itermax,mode,lscale);
```

The two snippets show the primary interfaces (in Matlab) to 2D Field Alignment routines, namely, *FA2DImNoH* and *FA2DC4V2*. Please see other examples for additional use cases. The interface *FA2DImNoH* is one we anticipate to be more useful and it has a C/Fortran interface that is implemented in *example2.c* and *example2.f90*.

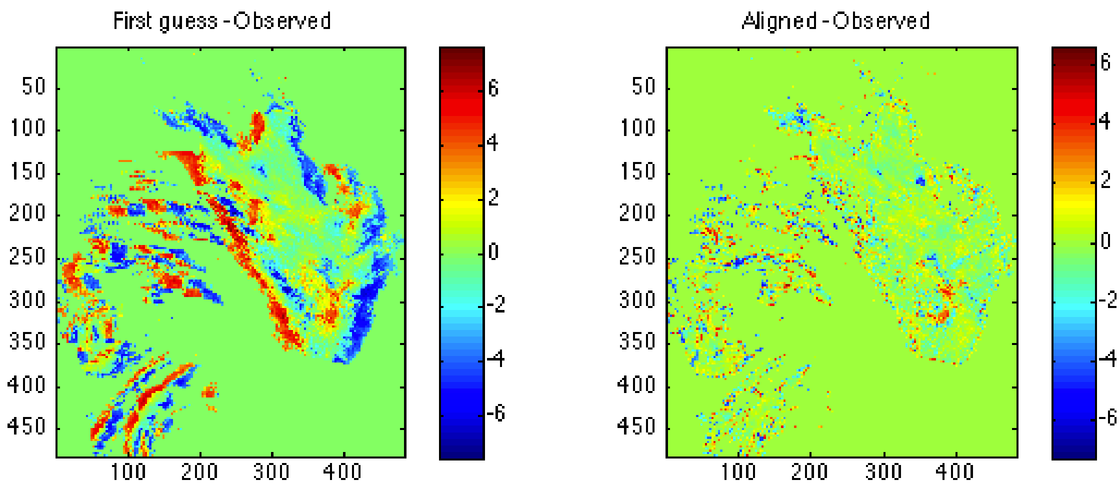

Figure 1: The error between two Radar reflectivity fields before and after applying Field Alignment. The dissipation of motion related error is significant.

Figure 1 shows a portion of the output from *example1*. Field Alignment was used to recover the apparent motion between a pair of radar images and the error after alignment shows substantial reduction from the error before, suggesting of course that most of the error is due to motion and that it has been well recovered.

Figure 2 similarly shows the reduction in error between two time separated (24h) GFS fields, where much of the error is explained in deformation space and the residual shows localized growth or decay.

FA can also be applied to multivariate fields as Figure 3. Here, Field Alignment is also used in a "spot alignment" mode, i.e. restricted within a region containing a vortex. The extended

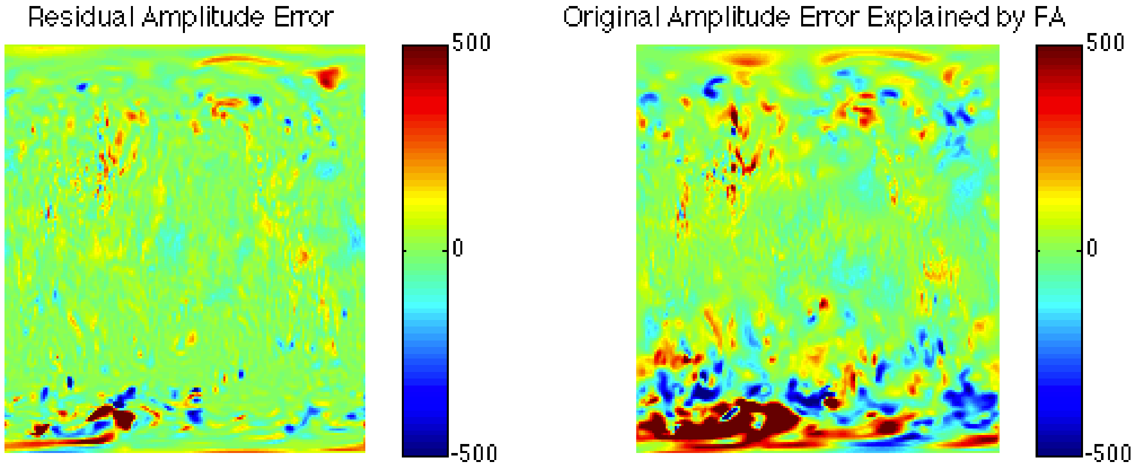

Figure 2: The error between two GFS fields before and after applying Field Alignment. The dissipation of motion related error is substantial.

domain beyond the local area is left undisturbed. The fields are multivariate, including pressure and velocity extracted from two GFS model outputs and with sparse measurements. Significant reduction in error is observed. FA can also be used in ensemble mode by aligning each ensemble member in-turn or in parallel, please see *example4.m*.

Last, but not the least, we turn to a 1D example implementing a new solver (*example1.m* in FA1D). In Figure 4, the alignment of a Gaussian bump is shown with sparse noisy observations that is 6σ away from a first guess that a local solution will not be able to recover. The interesting element of this solution is that the constraints on motion encoded in Λ are interpreted as a distribution in a way that allows for stochastic optimization to conveniently proceed. Other solvers using MCMC techniques have also been implemented and available upon request.

We hope that the FAST release will allow the user to undertake other applications and extend it in interesting ways. We would be delighted to help support such efforts to the degree possible.

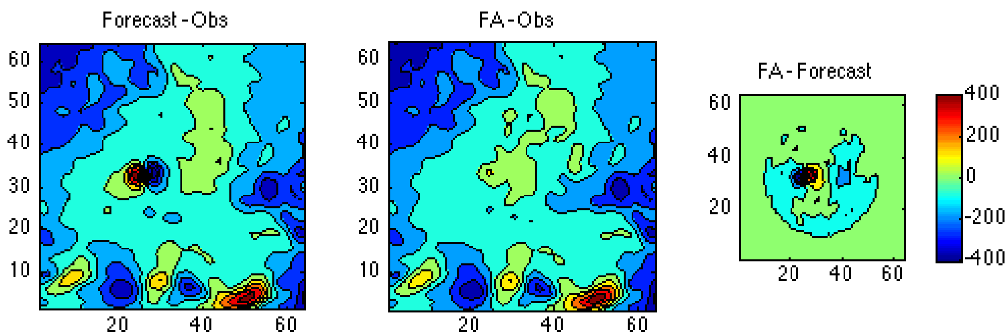

Figure 3: A multivariate alignment example, here only the pressure field is shown.

5 Upgrades, Updates and Limitations

The user can access a feed of updates and upgrades directly from MIT at http://tinyurl.com/fastmit. From MIT, we will automatically produce FAST upgrades twice each year and push it to Zeus. The next release will be sometime in January and will contain the following upgrades:

- Hierarchical Field Alignment: The hierarchical algorithm enables dealing with very large displacements and large fields (e.g. 1Kx 1K or greater). To be sure, the current approach will still be able to run them, but small-scale motions, more than frequency of 0.1cycles/pixel

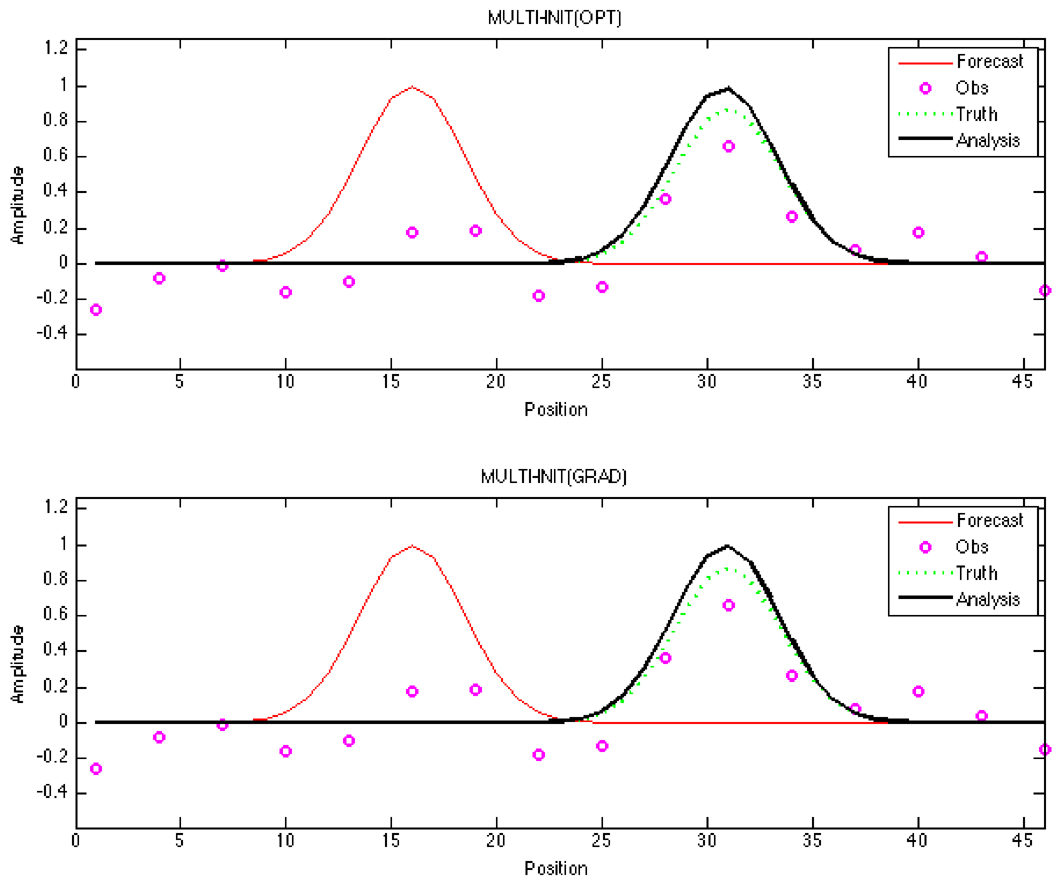

Figure 4: A one-dimensional example with a non-local solution. Above, stochastic BFGS, below stochastic gradient-descent.

may not be resolved.

- Localization: The localization is currently in the spectral domain, but for most applications one can trade the spectral and spatial domains for effectiveness. This will also be available in the next version.

- Stochastic Optimization: The stochastic optimization framework will be extended further in 2D.

References

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

[2] S. Ravela. Quantifying uncertainty for coherent structures. *Procedia Computer Science*, 9:1187–1196, 2012.

[3] S. Ravela. Storm data assimilation by appearance and geometry. *Proc. ICCV (submitted)*, 2013.

[4] S. Ravela, K. Emanuel, and D. McLaughlin. Data assimilation by field alignment. *Roger F. Daley Memorial Symposium*, 2003.

[5] Sai Ravela, Kerry Emanuel, and Dennis McLaughlin. Data assimilation by field alignment. *Physica D: Nonlinear Phenomena*, 230(1-2):127 – 145, 2007.

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

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