

Data assimilation in the solar wind: challenges and first results

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Lang, M., Browne, P., Van Leeuwen, P. J. and Owens, M.
ORCID: <https://orcid.org/0000-0003-2061-2453> (2017) Data
assimilation in the solar wind: challenges and first results.
Space Weather, 15 (11). pp. 1490-1510. ISSN 1542-7390 doi:
10.1002/2017SW001681 Available at
<https://centaur.reading.ac.uk/72986/>

It is advisable to refer to the publisher's version if you intend to cite from the
work. See [Guidance on citing](#).

Published version at: <http://dx.doi.org/10.1002/2017SW001681>

To link to this article DOI: <http://dx.doi.org/10.1002/2017SW001681>

Publisher: American Geophysical Union

All outputs in CentAUR are protected by Intellectual Property Rights law,
including copyright law. Copyright and IPR is retained by the creators or other
copyright holders. Terms and conditions for use of this material are defined in
the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

RESEARCH ARTICLE

10.1002/2017SW001681

Key Points:

- The paper describes the first data assimilation (DA) experiments with a solar wind model
- These experiments show that DA can improve forecasts of nonmagnetic variables near the observation, which are advected by the solar wind
- We discuss challenges and obstacles inherent to DA in the solar wind and propose potential solutions as directions for future research

Correspondence to:

M. Lang,
matthew.lang@lscce.ipsl.fr

Citation:

Lang, M. S., Browne, P. A., van Leeuwen, P. J., & Owens, M. (2017). Data assimilation in the solar wind: Challenges and first results. *Space Weather*, 15, 1490–1510. <https://doi.org/10.1002/2017SW001681>

Received 15 JUN 2017

Accepted 26 SEP 2017

Accepted article online 3 OCT 2017

Published online 16 NOV 2017

Data Assimilation in the Solar Wind: Challenges and First Results

Matthew Lang¹ , Philip Browne^{2,3}, Peter Jan van Leeuwen^{2,4}, and Mathew Owens² 
¹Le Laboratoire des Sciences du Climat et de l'Environnement, CEA-CNRS-UVSQ, Gif-Sur-Yvette, France, ²Department of Meteorology, University of Reading, Reading, UK, ³Currently at European Centre for Medium-Range Weather Forecasts, Reading, UK, ⁴NCEO, University of Reading, Reading, UK

Abstract Data assimilation (DA) is used extensively in numerical weather prediction (NWP) to improve forecast skill. Indeed, improvements in forecast skill in NWP models over the past 30 years have directly coincided with improvements in DA schemes. At present, due to data availability and technical challenges, DA is underused in space weather applications, particularly for solar wind prediction. This paper investigates the potential of advanced DA methods currently used in operational NWP centers to improve solar wind prediction. To develop the technical capability, as well as quantify the potential benefit, twin experiments are conducted to assess the performance of the Local Ensemble Transform Kalman Filter (LETKF) in the solar wind model ENLIL. Boundary conditions are provided by the Wang-Sheeley-Arge coronal model and synthetic observations of density, temperature, and momentum generated every 4.5 h at 0.6 AU. While in situ spacecraft observations are unlikely to be routinely available at 0.6 AU, these techniques can be applied to remote sensing of the solar wind, such as with Heliospheric Imagers or interplanetary scintillation. The LETKF can be seen to improve the state at the observation location and advect that improvement toward the Earth, leading to an improvement in forecast skill in near-Earth space for both the observed and unobserved variables. However, sharp gradients caused by the analysis of a single observation in space resulted in artificial wavelike structures being advected toward Earth. This paper is the first attempt to apply DA to solar wind prediction and provides the first in-depth analysis of the challenges and potential solutions.

1. Introduction

Variability in the solar magnetic field over minutes, hours, and days results in near-Earth solar wind conditions which can adversely affect space- and ground-based technologies (Cannon et al., 2013). The most extreme “space weather” is driven by coronal mass ejections (CMEs), large episodic eruptions of plasma and magnetic field from the Sun that travel through the solar wind, disturbing the Earth’s magnetic field as part of a geomagnetic storm (Gosling, 1993). As such, Earth-directed CMEs pose a threat to electrical and communication systems which have become a huge part of modern-day life, with an estimated potential economic impact of up to \$2 trillion in the first year after an extreme storm (Board, 2008). In addition, the radiation hazard of energetic charged particles associated with the solar wind and by CMEs is a health threat to humans at high altitude, such as aircrew over the poles and astronauts, both in low Earth orbit (e.g., on the International Space Station), but particularly on interplanetary missions, when the protection of the Earth’s magnetic field is removed (Cannon et al., 2013). The UK has responded to the danger of a major space weather event, such as a large CME, by adding it to the National Risk Register as one of the largest threats to modern society. The UK Met Office (UKMO) and U.S. Space Weather Prediction Center (SWPC) are working together to prepare for, and mitigate against, the possible impacts of a large space weather event. Thus, forecasting solar wind conditions in near-Earth space is a high priority.

The state-of-the-art method for forecasting near-Earth solar wind conditions is through coupled coronal and heliospheric models, with boundary conditions ultimately set by photospheric magnetic field observations (e.g., see Figure 1). As a complete photospheric magnetic field map takes one solar rotation to construct (27 days from Earth’s point of view), this approach enables the quasi steady state solar wind to be predicted. Time-dependent phenomena, such as CMEs, are then incorporated through ad hoc perturbation of heliospheric boundary conditions. This typically entails a “cone model” which inserts an overpressurized density

©2017. The Authors.

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.




Figure 1. The current state-of-the-art approach to operational numerical space weather forecasting. (top row) The coupled “chain” of MHD models from the solar photosphere to terrestrial thermosphere, all ultimately initiated using photospheric magnetic field observations. (bottom row) Current ad hoc solutions to key missing elements to this chain, namely, the time-dependent coronal structures, which are incorporated through “cone model” representation of CMEs, and energetic particles in the heliosphere and magnetosphere, incorporated through stand-alone kinetic codes. Adapted from Owens et al. (2014).

perturbation with the speed, direction, and angular width set by coronagraph observations of the CME in question (e.g., see Parsons et al., 2011; Xie et al., 2004). No additional observational constraints are imposed on the forecast between the top of the solar corona (20 solar radii) and near-Earth space (215 solar radii). For operational solar wind forecasting, UKMO and SWPC use ENLIL (Odstrcil, 2003; Odstrcil et al., 2004; Odstrcil & Pizzo, 1999; Parsons et al., 2011), a 3-D magnetohydrodynamic (MHD) model.

In this paper, we investigate the use of an advanced data assimilation method with the ENLIL model for potential improvement of space weather forecasting. To the authors’ knowledge, this is the first study to apply data assimilation methods to the solar wind. We use the EMPIRE (Browne & Wilson, 2015) data assimilation framework. In the next section the solar wind is described. Then the basic ingredients are described in succession, a short introduction to data assimilation, the model ENLIL, and the EMPIRE data assimilation framework. This is followed by the setup of the data assimilation experiments, the results, and a concluding section.

2. The Solar Wind

The solar wind is a plasma composed primarily of electrons and protons, with a small contribution from alpha particles and other minor species. It continually flows almost completely radially away from the top of the Sun’s hot corona (Parker, 1958), generating the heliosphere. The mechanism(s) by which the corona is heated, however, are still debated (e.g., De Moortel and Browning, 2015), as are the precise mechanism(s) by which the coronal plasma is confined and released (Cranmer, 2008; McComas et al., 2007). The solar wind drags the coronal magnetic field outward generating the heliospheric magnetic field (HMF) (Owens & Forsyth, 2013), which magnetically couples the Sun and planets and modulates the flux of galactic cosmic rays in the inner heliosphere. The solar wind plasma is accelerated to speeds greater than the characteristic wave speeds (i.e., the local Alfvén and fast magnetosonic wave speeds) within 0.1 AU (where 1 AU $\approx 1.50 \times 10^8$ km is the astronomical unit, defined as the average distance between the Earth and Sun), meaning that no information from the solar wind beyond this distance can propagate back toward the Sun. Typical properties of the solar wind at 1 AU are summarized in Table 1.

Table 1
Observed Properties of the Solar Wind Near the Orbit of the Earth (1 AU)

Properties	Values
Proton number density	$6.6 \times 10^{-6} \text{ m}^{-3}$
Electron number density	$7.1 \times 10^{-6} \text{ m}^{-3}$
Alpha particle number density	$2.5 \times 10^{-7} \text{ m}^{-3}$
Flow speed (nearly radial)	$4.5 \times 10^5 \text{ ms}^{-1}$
Proton temperature	$1.2 \times 10^5 \text{ K}$
Electron temperature	$1.4 \times 10^5 \text{ K}$
Magnetic field (induction)	$7 \times 10^{-9} \text{ T}$
Typical time for solar wind to flow from Corona to 1 AU	4 days = $3.5 \times 10^5 \text{ s}$

Note. Source: Kivelson and Russell (1995).

The ambient solar wind has two distinct components, typically referred to as the “fast” ($\sim 750 \text{ km s}^{-1}$) and the “slow” solar wind ($\sim 400 \text{ km s}^{-1}$), although their differences are not limited to only their speed (and extend to their formation mechanism at the Sun, for example). Fast solar wind flows outward along “open” magnetic field lines associated with coronal holes (Hassler et al., 1999), which are largely confined to the polar regions at times around sunspot minimum. The slow solar wind emanates from, or regions close to, closed coronal loops, which are confined to equatorial regions at sunspot minimum. During solar maximum, the coronal magnetic field is far more dynamic, with a much weaker dipole component. Consequently, the slow solar wind is observed at a much greater range of solar latitudes.

In situ spacecraft observations of the solar wind provide direct measurements of the solar wind plasma and HMF, which is directly rebateable to state vectors of the ENLIL simulation. However, only sparse spatial coverage of the global heliosphere is possible. Temporally, there has been near-complete data coverage in near-Earth space since 1996. Close to the L1 Lagrangian point between the Sun and the Earth (at heliocentric distance of $\sim 0.99 \text{ AU}$), Advanced Composition Explorer (ACE) (Stone et al., 1998) and Wind (Gloeckler et al., 1995) spacecraft have provided measurements since 1996, with the DSCOVR mission (Burt & Smith, 2012; Leslie & Cole, 2016) taking over as the standard observatory in June 2015. Away from near-Earth space, the two STEREO spacecraft (Davis et al., 2009; Kaiser et al., 2008) observe the solar wind from the elliptic plane, drifting 22° further ahead and behind the Earth each year, and the Helios spacecraft (Jackson, 1985a, 1985b), which explored the inner heliosphere to 0.29 AU . These data are all freely available online (e.g., at <https://nssdc.gsfc.nasa.gov/space/>).

From the available data it is difficult to forecast the near-Earth solar wind directly. Photospheric extrapolation enables a quasi-synoptic estimate of the near-Sun solar wind conditions, but it is highly indirect and therefore subject to large uncertainties. The sparse in situ observations are more reliable but sample only a small region of space; the L1 point is too close to Earth to provide a useful forecast lead time, and there are only extremely infrequently “upwind” monitors on the Earth-Sun line (e.g., the Helios spacecraft for very short intervals during the late 1970s and early 1980s). Heliospheric Imager (HI) type instruments, such as those on board STEREO (Eyles et al., 2009), and Interplanetary Scintillation (IPS) measurements (Manoharan & Ananthakrishnan, 1990) enable greater spatial sampling solar wind for limited properties (primarily density but also solar wind speed). Interpretation of these observations requires deconvolution of geometrical and path-integrated effects; thus, there is large uncertainty in, for example, solar wind speed, at a given point in the heliosphere. This is where data assimilation can be of assistance because it incorporates a full dynamical model of the solar wind, to some extent interpolating the observations to obtain a more complete picture, and allows for observational and model uncertainty. Thus, it provides a natural starting point for prediction. Here we explore how these observations can be brought together. We use a solar wind model (ENLIL) initiated using the near-Sun solar wind conditions from photospheric magnetic field observations and assimilate solar wind conditions at a single point, located 0.6 AU along the Sun-Earth line. In the remainder of the study, we consider these solar wind observations to be synthetic in situ spacecraft observations, but they could equally represent HI or IPS observations.

3. Short Introduction to Data Assimilation

Data assimilation combines prior information of a system encoded in a numerical model with new observational information to obtain a better description of the evolution of that system, including uncertainty estimates. The information in each component is represented by a probability density function (pdf). Data assimilation is based on Bayes theorem, which tells us that the prior pdf of the model has to be multiplied with the pdf of the observations to obtain the posterior pdf of the model given the observations:

$$p(x|y) = \frac{p(y|x)}{p(y)}p(x), \quad (1)$$

in which x denotes the model and y the observations, and the functions $p(\cdot)$ denote the different pdfs, distinguished by their arguments. The model x can be the state of the model at a specific time or a model trajectory over a certain time window. Similarly, the observations y can denote observations at a specific time or over a set of observations over a certain time window. If x is the model state at a specific time and y only contains observations at or before that time, it is referred to as filtering, while smoothing refers to the situation that y contains observations that occur at a later time step than x , and x can be a state or a trajectory over time. If x and y refer to different times, the Bayesian framework needs the joint pdf of the state, x , and the state at the time of the observation to bring the observation information to the state, x , at the time of interest. Details can be found in the literature and at the end of this section.

The state of the model is described by a state vector $\mathbf{x} \in \mathbb{R}^{N_x}$, which contains the values of the quantities of interest at all grid points, where N_x is the dimension of the state. For ENLIL, the state vector contains the three components of the magnetic field vector, the three components of the solar wind plasma momentum vector, the plasma temperature, the plasma density, and the cloud and polarity tracers, at each grid point of the model.

The evolution of the model is described by the model equation

$$\mathbf{x}_{i+1} = f_i(\mathbf{x}_i) + \boldsymbol{\eta}_i, \quad (2)$$

where $\mathbf{x}_i$ is state vector at time step i , f_i represents the pure model that incorporates our understanding of the physics of the system, and $\boldsymbol{\eta}_i$ is an N_x -dimensional stochastic term that represents the error in the model equations, resulting from discretization errors, missing physics, and inaccurate boundary conditions. In the case of solar wind forecasting, the inaccurate boundary conditions are expected to be the largest single factor. In data assimilation the statistics of the model errors is assumed to be known, for example, Gaussian with zero mean and prescribed error covariance. This, of course, is an idealization, and estimating these statistics is becoming an active field of research in, for example, weather prediction, where, traditionally, these model errors have been ignored for computational reasons.

In the geosciences we never know the initial state vector, and in data assimilation it is typically assumed to be a random perturbation from the true state:

$$\mathbf{x}^b = \mathbf{x}_0^t + \boldsymbol{\xi}_0, \quad (3)$$

where $\mathbf{x}^b \in \mathbb{R}^{N_x}$ is the initial state, $\mathbf{x}_0^t \in \mathbb{R}^{N_x}$ is a discretization of the true state at time 0, and $\boldsymbol{\xi}_0 \in \mathbb{R}^{N_x}$ represents the random error in the initial state. Since the truth is not known, the initial condition is either based on a previous forecast or is carefully generated using all physical knowledge at hand.

The observations are measurements of the true system but contain measurement errors. They also contain errors arising from the fact that model and observations tend to represent reality differently; for example, they have different spatial resolution. Quantifying these so-called representation errors is again an active area of research; see, for example, Hodyss and Nichols (2015) and van Leeuwen (2015). The relation between observations and state of the system is written as

$$\mathbf{y}_i = \mathcal{H}_i(\mathbf{x}_i^t) + \boldsymbol{\epsilon}_i, \quad (4)$$

where $\mathcal{H}_i : \mathbb{R}^{N_x} \rightarrow \mathbb{R}^{N_y}$, the observation operator, maps the state into observation space and N_y is the dimension of the observations. The observation error, containing both instrument and representation error, is given by $\boldsymbol{\epsilon}_i$, the statistics of which we assume to be known.

Data assimilation can be used for state estimation, as described above, but it is also used for parameter estimation (see, e.g., Evensen et al., 1998; Smith et al., 2009) and has recently being applied to the estimation of nonlinear parameterizations (Lang et al., 2016). In this paper, however, we shall focus upon state estimation.

It is important to realize that there is a fundamental problem in data assimilation for the geosciences that has to do with the size of the problems involved. Suppose we want to store the prior pdf of a 100-dimensional system, which is a relatively small system. Since that pdf can have any shape, we would have to rely on histogram representations. Assume that we use 10 frequency bins for each variable, then we need to store of the order of 10^{100} numbers. Our present-day supercomputers can store a lot of numbers, but this is completely out of the question. By comparison, the number of atoms in the observable universe is estimated to be of the order of 10^{80} , so the data assimilation problem is larger than the observable universe! This estimate is very conservative; the dimension of the state vector for the low-resolution ENLIL simulations performed in this study is 3,888,000. This means that one has to make approximations to the full Bayesian solution, and a large part of data assimilation research is focussed on finding the best approximation for specific problems. This has resulted in a number of different data assimilation methods, and the following gives a very brief overview of what is currently used in the geosciences. More detailed information can be found in recent text books like Nakamura and Potthast (2015), Reich and Cotter (2015), and van Leeuwen et al. (2015).

3.1. Variational Methods

While variational methods are not used in this paper, for reasons discussed below, it is useful to give an overview of the approach.

In variational methods one tries to find the mode of the posterior pdf, which is the most probable state of the model given the observations. The mode is defined as

$$\hat{\mathbf{x}} = \operatorname{argmax}_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}). \quad (5)$$

Writing $p(\mathbf{x}|\mathbf{y}) = \exp[-\mathcal{J}(\mathbf{x})]$, this amounts to minimizing the cost function $\mathcal{J}$, which is done by setting the gradient of that function to 0 and deriving the Euler-Lagrange equations, and solving these with an efficient gradient descent algorithm, such as conjugate gradient. This involves generating the tangent linear model and its transpose, the adjoint, of the nonlinear model code, which can be a formidable task.

As an example, when we assume Gaussian errors statistics with zero means and covariances; $\mathbf{R}_k$ for observational errors, $\mathbf{B}$ for initial condition errors, and $\mathbf{Q}_i$ for model errors, we find this cost function:

$$\begin{aligned} \mathcal{J}(\mathbf{X}_{N_t}) = & \frac{1}{2} (\mathbf{x}_0 - \mathbf{x}^b)^T \mathbf{B}^{-1} (\mathbf{x}_0 - \mathbf{x}^b) + \frac{1}{2} \sum_{k=1}^d \left[(H_k(\mathbf{x}_{i_k}) - \mathbf{y}_k)^T \mathbf{R}_k^{-1} (H_k(\mathbf{x}_{i_k}) - \mathbf{y}_k) \right] \\ & + \frac{1}{2} \sum_{i=0}^{N-1} \left[(\mathbf{x}_{i+1} - \mathbf{f}(\mathbf{x}_i))^T \mathbf{Q}_i^{-1} (\mathbf{x}_{i+1} - \mathbf{f}(\mathbf{x}_i)) \right], \end{aligned} \quad (6)$$

where $\mathbf{X}_{N_t} = (\mathbf{x}_0, \dots, \mathbf{x}_{N_t})^T$ is the vector of the states at all time steps, N_t , and i_k denotes the time steps that we have observations in the time window. Minimizing this cost function is the so-called weak-constraint 4DVar problem; 4D because it contains space and time, and weak constraint because it allows for errors in the model equations. When these are ignored, the last term is absent and we obtain strong-constraint 4DVar.

For numerical weather prediction, and also for space weather applications, the covariance matrices involved are huge, for example, 10^{18} for weather prediction, so they cannot be stored explicitly and are coded as operators working on input vectors. Furthermore, for these very high-dimensional optimization problems, preconditioning is essential. For strong-constraint 4DVar the matrix $\mathbf{B}$ is often used, and it is still unclear what the best preconditioning is for the weak-constraint 4DVar.

3.2. Ensemble Kalman Filters

While variational methods are very powerful, they have a few major drawbacks. The first one is that coding up the tangent linear and adjoint of the nonlinear forward model is a serious exercise, and for a complex model like ENLIL, this could easily take up a person-year or more. Furthermore, when the problem is highly nonlinear the posterior pdf can have multiple modes, and finding the global mode is not trivial. Finally, when the initial state errors are not Gaussian, standard optimization methods cannot be applied and very little experience is available in the community.

An alternative is sequential methods that forecast the initial state, $\mathbf{x}^b$, and its uncertainty to the first observation time step, i_0 , using the numerical model. Bayes theorem is used to update the pdf, and a new forecast is made to the next observation. The majority of sequential methods are based upon the Kalman filter (Kalman, 1960; Kalman & Bucy, 1961). Kalman filter-based methods search for the minimal variance state, which is the mean of the posterior pdf. They assume that all pdfs are Gaussian, so only the mean and the covariance are needed to describe all pdfs involved. The main advantage of the approach is that a sequential method is typically simpler to implement and does not require the computation of the adjoint of the forward model. The Kalman filter also explicitly evolves the forecast state covariance matrix, $\mathbf{P}^f$, through the assimilation window. However, this requirement also renders the Kalman filter impractical for use in high-dimensional systems, as it may not be possible to explicitly store the forecast state error covariance matrix, let alone evolve it forward in time using the numerical model.

To solve this issue, the pdf is represented by a number of samples or ensemble members, and these are propagated to observation times using the full nonlinear model equations. At the observation time, the Kalman filter equations for updating mean and covariance are used to define the posterior pdf, using the ensemble members to provide prior estimates for them. By using the fact that the number of ensemble members is typically much smaller than the system dimension, very efficient updating schemes have been developed. This has resulted in a large class of so-called ensemble Kalman filters (EnKFs), a set of Monte Carlo-based sequential DA methods.

There are two main issues that hinder application in high-dimensional geophysical systems. First, due to the small ensemble size M , typically in the range 10–500, which is much smaller than the system dimension, N_x , the ensemble covariance matrix has rank $M-1$, which is much smaller than full rank. It spans only the directions of the ensemble perturbations from the ensemble mean, and hence, the covariance is estimated from below. To compensate for this underestimation (and typically also to compensate for an underrepresentation of the model errors), covariance inflation is applied. There are several ways to do this, and the most widely used method is simply multiplying the ensemble covariances by a factor $1 + \rho > 1$, as follows:

$$\tilde{\mathbf{P}}_{\text{infl}}^f = (1 + \rho) \mathbf{P}^f, \quad (7)$$

in which $\mathbf{P}^f$ is the forecast ensemble at the observation time and $\tilde{\mathbf{P}}_{\text{infl}}^f$ represents the inflated forecast error covariance matrix. This multiplication has the effect of increasing the spread in the forecast ensemble to better represent the true errors in the model. However, the uncertainty is only increased in the directions already covered by the ensemble. Other methods for covariance inflation have been developed; see for example, Anderson (2007) for an overview.

The second issue is that the covariances are estimated directly from a rather small ensemble, so especially small covariances are prone to sampling noise. Since covariances are expected to be small between distant grid points, so-called covariance localization can be applied, in which the sample covariance is multiplied, via a Schur product, with a localization matrix that tapers off quickly with grid point distance (Houtekamer & Mitchell, 2001). This procedure ensures that the spurious correlations between points are reduced and hence unrealistic covariances do not affect the analysis of the state. Another advantage of localization is that it makes the ensemble covariance matrix more diagonal, resulting in a localized ensemble covariance matrix that is, typically, of higher rank.

However, as mentioned before, the ensemble covariance is never calculated explicitly as it would not fit in memory leading to a different localization strategy called observation localization. In this method each grid point is treated separately by selecting a localization radius, and only observations within that radius are allowed to update this grid point. Within the localization radius the observation covariance matrix is multiplied with a factor increasing with the distance of the observation to the grid point, so that observations further away have a larger error and hence have less effect on the update of that grid point. This ensures a smooth spatial update. It is important to realize that inflation and localization are essential ingredients in ensemble Kalman filtering.

In this study, we shall use the LETKF, as developed by (Hunt et al., 2007), which uses observation localization. The LETKF is one of the most efficient and accurate methods currently used in operational weather prediction centers, such as the Japanese Meteorological Agency (JMA) (Miyoshi et al., 2010) and the National Center for

Environmental Prediction (NCEP) (Szunyogh et al., 2008). Appendix A contains a description of the full LETKF algorithm, showing how the whole algorithm works very efficiently in the low-dimensional ensemble space.

3.3. Other Data Assimilation Algorithms

Recent years have seen a surge of other Monte Carlo-based methods for geophysical problems, mainly particle filters. Their main advantage over the methods discussed above is that they are fully nonlinear. However, until recently these methods were thought to be too inefficient to be useful in high-dimensional systems. But by exploring either the proposal density freedom or localization, efficient algorithms have been developed that are very promising for high-dimensional geophysical problems (see recent overviews in Reich & Cotter, 2015; van Leeuwen et al., 2015).

Another active area of research is in hybrid methods, where different data assimilation methods are combined, exploiting the strengths of each. Examples are ensemble smoothers like 4DEnsVar, in which an ensemble of model runs is used to generate space-time covariances that alleviate the need for tangent linear and adjoint models in a 4DVar algorithm. We will not discuss these developments here but refer to recent papers like Fairbairn et al. (2014), Goodliff et al. (2017) and Amezcua et al. (2017, and references therein).

4. The ENLIL Heliospheric Solar Wind Model

The dynamics of the solar wind differ from the typical dynamics of the atmosphere, as the solar wind is strongly driven by the conditions at the top of the corona (i.e., the inner boundary condition to the ENLIL model), whereas typical atmospheric systems are highly chaotic and so are very sensitive to small perturbations. This has direct consequences for the data assimilation, as we will see later.

Over time and length scales greater than the electron and ion gyromotions, which are many orders of magnitude below typical solar wind model resolutions, the large scale behavior of the solar wind plasma can be well represented as a magnetized fluid, the magnetohydrodynamic (MHD) approximation. Under such conditions, the fluid has negligible resistivity and so can be treated as a perfect conductor. This means that, through Lenz's law, the motion of the plasma and the magnetic field is "frozen" together; as long as resistivity remains negligible, the magnetic field within the plasma of the solar wind will move with the velocity of the plasma (e.g., Kivelson & Russell, 1995). Ideal MHD, the simplest approximation of MHD, therefore reduces to the continuity equation, the Cauchy momentum equation, Ampère's law, and a temperature evolution equation. Thus, ideal MHD explicitly represents the conservation of mass, momentum, total energy, and induction of the magnetic field, as well as the effects of the magnetic field via magnetic field pressure and tension of magnetic field lines (Odstrcil, 2004) (the force exerted by the curvature in a magnetic field line as it tries to "straighten" out).

The ENLIL model is a 3-D numerical model that is used operationally at the UK Met Office and NOAA's Space Weather Prediction Centre in combination with coronal models and magnetospheric models, as described in Figure 1. It is based upon the ideal magnetohydrodynamics (MHD) equations, with two additional continuity equations for the "Cloud Tracer," a passive tracer that traces the material from a cone model CME, and the "Polarity Tracer," a passive tracer for the polarity of the HMF (see Odstrcil and Pizzo, 1999 for more details). The latter is required as the HMF is treated as a unipolar field at run time, in order to avoid the HMF numerically diffusing away close to the polarity inversion (i.e., the heliospheric current sheet). It is then returned to its bipolar state postrun. Numerical calculations are performed in a spherical coordinate system on a domain decomposition approach to divide the 3-D computational domain into smaller radial slabs for processing on parallel systems, via MPI (Message Passing Interface) (Gropp et al., 1996). Each processor then solves the MHD equations on the respective radial slabs and boundary data for each slab is exchanged via MPI calls.

The inner boundary of the ENLIL model is specified at ~ 0.1 AU, outside the point at which the solar wind becomes supersonic. Thus, there is no sunward propagation of information through the inner boundary, simplifying the numerical computation. To specify the inner boundary conditions, the Wang-Sheeley-Arge (WSA) model (Arge & Pizzo, 2000; Wang & Sheeley Jr. 1992), a semiempirical model of the corona, is typically used. The WSA model takes observations of photospheric magnetic field as input and extrapolates the field through the corona, typically to a source surface of around 2.5 solar radii, using the potential-field source surface approximation. The solar wind velocity, proton density, and temperature can be derived using empirical relationships to the coronal magnetic field, assuming a constant mass flux (see Lopez, 1987; Riley et al., 2015 for more details). These properties are radially mapped to 0.1 AU (ENLIL's inner boundary). All of these structures




Figure 2. Schematic of the EMPIRE data assimilation framework.

codes have been optimized to run in parallel, via MPI commands, such that large ensembles can be processed efficiently and effectively. At observation time, the model passes a state vector to EMPIRE, which performs the data assimilation independent of the model and then passes the analyzed state vector back to the model and waits until the next observation time step.

There are other data assimilation libraries available such as the Data Assimilation Research Testbed (DART) (Anderson et al., 2009) and Parallel Data Assimilation Framework (PDAF) (Nerger et al., 2005). While these data assimilation systems are highly optimized, their main strength is that they are focussed on academic users. This means that minimal coding is needed to couple any numerical model to these systems, exploring the fact that most ensemble-based data assimilation methods are not dependent on how the model works. Furthermore, once coupled it is extremely easy to switch to data assimilation method, allowing fast comparisons and a fast way to choose the best method for the problem at hand. Additionally, EMPIRE and PDAF have the additional advantage that it is not necessary for the model to write to disk every assimilation step, due to the efficient use of MPI and/or subroutine calls, meaning more efficient computations. Moreover, all these systems are community code with active user support.

As the DA codes are completely general, it is necessary to specify the observation time steps, the observation operator, the error covariance matrices, and distances between two points in the state vector (for the LETKF) only. This makes the implementation of EMPIRE into the numerical model relatively simple in comparison to DART, which requires the model code to be adjusted to a higher degree than the simple MPI command required for the implementation of EMPIRE (Browne & Wilson, 2015).

6. Numerical Experiments

Since this paper describes initial tests of a data assimilation experiment, we perform the experiments in a controlled way. We start by recognizing that we have some uncertainty in the model state. With this uncertainty we are able to generate an ensemble which represents our prior pdf. Then, using the same technique as we use to generate each ensemble member, we generate one further model state which we refer to as the *truth* state. We use our numerical model to propagate this truth state to get a complete *truth trajectory*. We then take artificial observations from a fixed point in space from this truth trajectory. These “observations” are perturbed by measurement noise to mimic a real data assimilation experiment. The goal is to closely represent the posterior pdf, which represents the best estimate including an uncertainty estimate of the truth run, using the limited information from the observations and uncertain prior information on initial and boundary conditions. The ensemble generated is evolved using the numerical model in two separate runs, one with data assimilation performed and one without it to evaluate the performance of the data assimilation scheme used. This setup is called an identical twin experiment, and it is the first test for any data assimilation system.

6.1. Experimental Setup

Twin experiments have been performed using the LETKF to assimilate observations of the true state generated from an unknown initial condition.

The state vector for the solar wind field is defined as a vector of the density, temperature, momentum (radial, latitudinal, and longitudinal components), the magnetic field (radial, latitudinal, and longitudinal components), and cloud and polarity tracers at each point in the ENLIL domain. Writing this mathematically,

are rotating, with an azimuthal velocity equal to the Sun’s rotation speed along the inner boundary to create the inner boundary values for all time steps in the time domain.

5. The EMPIRE Data Assimilation System

To perform the data assimilation experiments efficiently, the EMPIRE data assimilation system (Browne & Wilson, 2015) is used. EMPIRE contains data assimilation codes that link to the model via a small number of MPI commands (Gropp et al., 1996), see Figure 2 for a visualization of this. The data assimilation methods in EMPIRE include some of the most advanced ensemble-based data assimilation methods such as the LETKF (which shall be used in this paper), the Bootstrap Particle Filter, the Equivalent Weights Particle Filter, the Implicit Equal Weights Particle Filter, and 4DEnVar. The

Table 2
Table Showing Setup for the Data Assimilation Experiments

DA method used	LETKF
No. of ensemble members	48
Length of spin-up	1,600 time steps (equivalent to 5.93 days)
No. of time steps after spin-up	1,350 (equivalent to 5 days)
State vector	$(\rho^T, \mathbf{T}^T, (\rho \mathbf{v}_r)^T, (\rho \mathbf{v}_\theta)^T, (\rho \mathbf{v}_\phi)^T, \mathbf{B}_r^T, \mathbf{B}_\theta^T, \mathbf{B}_\phi^T)^T$
Observations taken	All variables in state vector except magnetic field variables at spatial coordinate $(r, \phi) = (74, 0^\circ)$
Frequency of observations	Every 50 time steps (equivalent to every 250 min)
Initial error covariance	Specified by long-run ensemble snapshots
Observation error covariance	See Table 3
Model error covariance	$\mathbf{0}$
Localization radius	0.01 AU
Inflation factor	0

$$\mathbf{x} = (\rho^T, \mathbf{T}^T, (\rho \mathbf{v}_r)^T, (\rho \mathbf{v}_\theta)^T, (\rho \mathbf{v}_\phi)^T, \mathbf{B}_r^T, \mathbf{B}_\theta^T, \mathbf{B}_\phi^T, \rho_c^T, \rho_p^T)^T, \quad (8)$$

where the bold font denotes the relevant variable at all radial, latitudinal, and longitudinal points in the model's domain.

ENLIL is run with an inner boundary of 0.1 AU and outer boundary of 1.1 AU in a spherical grid with 144 radial points equally spread throughout the domain; 30 latitudinal points spread equally between 30° and 150° latitude (with 0° defined as the north pole of the Sun) and 90 longitudinal points spread equally between -90° and 90° with 0° defined as the line between the Sun and Earth. The time domain is 5 days with time steps every 320 s (plus an additional spin-up time of ~ 5.93 days).

6.1.1. Assimilation Parameters

In this section, the setup for the data assimilation experiments performed in this paper is described. The main parameters used for the assimilation are summarised in Table 2.

Boundary conditions for ENLIL were generated at 20 solar radii (≈ 0.1 AU) from outputs of the WSA model for a magnetogram from 1 April 2013. To generate the truth run, the ENLIL model was spun-up for 5.12×10^5 s (1,600 time steps or ≈ 5.93 days) in order for the inner boundary conditions to advect the solar wind structures through the whole model domain. These spun-up boundary conditions are then used as initial conditions for the data assimilation experiments. By spinning up the boundary conditions in this manner, as opposed to prescribing a set of initial conditions (from a long run, for example), we avoid initialization shocks occurring in the domain that may contaminate our results. The truth run started after spin-up and lasted for 5 days, using 1,350 model time steps.

Observations of the true state were taken every 50 time steps (approximately every 3 h). The observations were taken of density, temperature, and the momentum variables with an observation error covariance given by the diagonal matrix, $\mathbf{R} \in \mathbb{R}^{5 \times 5}$, where the variances along the diagonal are shown in Table 3. The variances are approximately 1–10% of typical variable values at the observation location. In the experiments shown here, observations are taken at a single point (to represent a spacecraft, or properties at a single point in

space determined from HI or IPS observations) in the ecliptic plane between the Earth and Sun. Observations are taken in the middle of the ENLIL domain, at approximately 0.6 AU. This was not done near the Earth (at 0.99 AU), as would be most realistic at the present, as this is close to ENLIL's outer boundary, and hence it could be not seen how the LETKF affected the state downwind. It is worth noting at this point, that after launch in 2019, ESA's Solar Orbiter and NASA's Solar Probe Plus will be taking observations of the solar wind at a radius of approximately 0.3 AU. While such observations will not be routinely available for forecasting, they will enable both more direct testing of the DA schemes and calibration of the near-Sun HI and IPS observations.

Table 3
Table Showing the Standard Deviations of the Errors in the Observations Taken of the Solar Wind

Variable observed	Standard deviation of observations
Density	$10^{-20} \text{ kg m}^{-3}$
Temperature	10^{-12} K
Radial momentum	$10^{-15} \text{ kg m}^{-2} \text{ s}^{-1}$
Latitudinal momentum	$10^{-17} \text{ kg m}^{-2} \text{ s}^{-1}$
Longitudinal momentum	$10^{-17} \text{ kg m}^{-2} \text{ s}^{-1}$

Two ensemble runs were used, a model-only ensemble run without data assimilation and an ensemble run where we use the LETKF. The model-only ensemble run represents the propagation of the prior pdf in time. By studying how the LETKF ensemble deviates from that pdf evolution, we can infer the impact of the data assimilation.

Our initial attempt at creating an ensemble was by adding random perturbations to the inner boundary condition generated from a diagonal initial error covariance matrix to an initial state generated using the WSA model boundary condition but found this to be ineffective (Lang, 2016). The spread ensemble was too small and simply increasing the magnitude of this diagonal covariance matrix causes the ENLIL model to become numerically unstable. Ideally, a covariance matrix that has proper correlations between all model variables should be used, but generating such a matrix is rather complicated and needs very long model runs. In NWP, the specification of such a covariance matrix can form a substantial part of the scientific effort to improve the state estimation.

Instead, here we use an ensemble of 48 members generated from snapshots of a long model run. Boundary conditions were drawn from 48 random WSA model outputs from the year 2015 (so they are independent from the true state). The random selection of the boundary conditions is done in order to ensure a wide variety of the Sun's behavior, both in terms of phasing of structures with solar rotation and changes in the patterns of the solar wind (e.g., slow/fast wind). Like the truth run, each ensemble member was spun-up for 1,600 time steps to avoid initialization shocks. This effectively swept out the model interior from the snapshots of the long model run and left us with internally balanced initial model states consistent with the boundary conditions.

Once the initial ensemble was specified from the boundary conditions, as described as above, the model-only ensemble was generated by propagating the 48 ensemble members through time to the end of the assimilation window of 5 days with 1,350 time steps, with no data assimilation. Using the same initial ensemble as the model-only ensemble, the LETKF was used to perform data assimilation experiments.

As mentioned previously in section 3, a localization function must be specified. In order to avoid upwind information propagation from the observation location back toward the Sun, an asymmetric radial distance-based localization function has been specified. Downwind of the observation location a truncated Gaussian function centered on the observation location was used with localization radius of 0.01 AU. Upwind of the observation the localization function is set to 0. Here this localization function multiplies the observation error covariance matrix, $\mathbf{R}^{-1}$. That is to say, the localization function only allows points within a radius of ~ 0.04 AU downwind of the observation to be updated and no points upwind of the observation to be influenced by the observation.

The experiments shown in this paper were run on 1,176 processors on ARCHER (Archer, <http://www.archer.ac.uk/>), the UK national supercomputing facility. Each processor has all variables at all latitudes and longitudes in the computational domain with eight radial points, the first and last radial values shared with the previous and next processor via MPI, respectively.

6.2. Diagnostics

In the following sections, for visualization purposes, the variables are multiplied by a factor of r^2 , the distance to the sun squared, to compensate for the r^2 decrease in density, momentum, and magnetic variables away from the Sun. This allows us to easily see structures in the state propagating further into the domain.

Results are analyzed over the ecliptic plane (i.e., the plane containing the Earth and the Sun) only. Results from the LETKF are sampled every 10 time steps. The radial points, $r = 2, 10, \dots, 138$, are sampled every eight grid cells (every ~ 0.0056 AU), the latitudinal coordinate is taken at only $\theta = 90^\circ$, and longitudinal coordinates, $\phi = -40^\circ, -30^\circ, \dots, 40^\circ$, are sampled every 10° around the Earth-Sun line.

The absolute error is also calculated at the observation point for the model-only ensemble and LETKF analysis ensemble mean. This absolute error, or RMSE, in the ensemble at the observation point is calculated by

$$|x_{r,\phi}^t - \overline{x_{r,\phi}}|, \quad (9)$$

where $x_{r,\phi}^t$ is the true state for a single variable at the point (r, ϕ) in the spatial domain, $\overline{x_{r,\phi}} = \frac{1}{M} \sum_{m=1}^M [x_{r,\phi}^m]$ is the ensemble mean for either the LETKF or model-only ensemble at a single point (r, ϕ) , and $x_{r,\phi}^m$ is the m th ensemble member at (r, ϕ) .




Figure 3. State at observation point, $(r, \phi) = (74, 0^\circ)$, for the different variables of the state vector for the LETKF and model-only ensemble run. The black line indicates the true state, the light and dark blue lines show the LETKF ensemble members and their mean, respectively, and the light and dark red lines show the model-only ensemble members and their mean.

The ensemble spread at each coordinate is also generated by calculating the standard deviation of the respective ensemble at each radial and longitudinal point, which is given by

$$\sqrt{\frac{1}{M-1} \sum_{m=1}^M \left[\left(x_{r,\phi}^m - \overline{x_{r,\phi}} \right)^2 \right]}. \quad (10)$$

The differences between the absolute errors for the LETKF and model-only ensemble are also calculated, at each point in the ecliptic plane sampled (for each (r, ϕ) point in equation (9), as

$$|x_{r,\phi}^t - \overline{x_{r,\phi}^{sto}}| - |x_{r,\phi}^t - \overline{x_{r,\phi}^a}|, \quad (11)$$

where $\overline{x_{r,\phi}^{sto}}$ represents the model-only ensemble mean and $\overline{x_{r,\phi}^a}$ represents the LETKF ensemble mean for a single variable at ENLIL coordinate (r, ϕ) on the ecliptic plane. Therefore, positive values (red regions on polar plots) indicate an improved LETKF ensemble, and negative values (blue regions on polar plots) indicate a poorer LETKF ensemble when compared to the model-only ensemble, which is comparable to no assimilation.

6.3. Results

To demonstrate the effectiveness and potential of the LETKF in assimilating solar wind observations, the density, radial momentum, and radial magnetic field are plotted at the observation location in Figure 3. Figure 3 shows the model-only ensemble and its mean, in red shades, and the LETKF analysis ensemble, in blue shades, compared to the truth in black, at the observation point, for each variable observed. It can be seen from Figure 3 that for the majority of time steps, the LETKF analysis state performs much better, for the density and radial momentum variables, at the observation location than the model-only ensemble when an observation is processed by the data assimilation scheme. At each observation time step, it can be seen that the LETKF trajectory is pulled toward the true state from the model-only ensemble. This results in reduced absolute errors in the LETKF ensemble when compared to the model-only ensemble. However, this improvement in the LETKF is quickly advected away from the observation location, as can be seen both in the ensemble mean and also in each ensemble member. This is to be expected as the model-only ensemble immediately replaces the updated ensemble in the time steps after assimilation due to the strong radial flow. Here one of the major differences with atmospheric weather data assimilation can be seen.

Observations of the radial magnetic field component are not currently used in the assimilation, to avoid violation of the $\nabla \cdot \mathbf{B} = 0$ condition. This is the reason why the pure and LETKF ensembles are almost identical.

Figure 4 shows the absolute error and ensemble spread for the pure and LETKF ensembles. The absolute errors for the pure and LETKF ensemble are shown by the blue and red lines, and the pure and LETKF ensemble spread is shown in dark blue and dark red, respectively. It is encouraging to see that the absolute error and the ensemble spread are of the same order of magnitude, showing that the ensemble is able to represent the uncertainty in the estimates well. The absolute error is more variable, but that is to be expected as it is a single realization of an error, while the ensemble spread is a statistical estimate. The nature of the assimilation updates on the spread and absolute error is consistent with the behavior of the ensemble members shown in Figure 3. Again, due to the strong solar wind, any model state update is advected away quickly from this grid point, leading to the strong variations in absolute error and spread. The LETKF ensemble spread is reduced to approximately the same value at each analysis time step and so does the absolute error, although less consistently. This is to be expected as the prior and the observation error variances have similar value just before each update step.

Figure 4 shows interesting behavior of the LETKF at some time instances when the absolute error becomes larger at an assimilation step. This is consistent with the behavior of the ensemble mean in Figure 3. The ensemble spread is always decreasing. The latter is consistent with what a Kalman filter-like update should do; the posterior covariance is always smaller than the prior covariance. For the RMSE we have to realize three things. First, it is a random realization of the error, so it can go up. Second, and perhaps more importantly, the updates are not univariate, but all variables are updated at the same time. It is well possible that a smaller RMSE for one variable is compensated by a larger absolute error in another, so that the total variance (defined, e.g., as the trace of the posterior absolute error matrix) is still decreasing. However, this is not what we expect given that we observe all variables at the same time. Third, it might be related to the fact that we do not update the magnetic field. This leads to slightly inconsistent updates, pushing the system out of balance. The model will react by an adjustment process that is typically wavelike. This will not be visible well at the observation location due to the strong solar wind but might be visible downwind. This seems to be the case, as discussed later.

Figure 5 shows the evolution of the difference of the absolute errors over the spatial domain, as calculated through equation (11). Red areas indicate that the LETKF has reduced the absolute error, while blue areas mean the opposite. The left column shows the density, the middle column the radial momentum, and the right column the radial magnetic field. The rows denote time instances 260 time steps apart, starting from 60 time steps into the assimilation run. This frequency is chosen to show the quality of the forecast produced by the LETKF ensemble 10, 20, ... time steps after the observation. All fields show positive impact of the LETKF, being advected toward Earth with the solar wind. In the density and radial momentum fields, we see a wavelike update instead of the expected steady update. It is clear that the model is adjusting itself from the assimilation updates by producing a wavelike feature which partly negates the positive impact of the assimilation. The impact on the radial magnetic field seems to be steady until a feature appears after about 1,100 time steps.




Figure 4. Absolute errors and ensemble spread at the observation point, $(r, \phi) = (74, 0^\circ)$, for the different variables of the state vector for the LETKF and stochastic ensemble run. The light and dark blue lines indicate the LETKF ensemble absolute error and ensemble spread, respectively, and the light and dark red lines indicate the model-only ensemble absolute error and ensemble spread, respectively.

To investigate this further, we produce Hovmöller plots along the radial direction for the three variables in Figure 6. These confirm the wavelike disturbance in the density and radial momentum fields. The waves are not completely regular but have an average period of about 500 time steps, corresponding to 1.85 days, with a wavelength of about 60 grid points, which corresponds to 0.45 AU. These features are likely to be the result of increased total pressure (the plasma density is increased without a compensating decrease in the magnetic field). Interestingly, this is equivalent to the current method for simulating CMEs within ENLIL, that is, the insertion of an overpressured density perturbation at the inner boundary.

7. Data Assimilation Challenges for Space Weather Models

A major challenge facing the implementation of variational methods in space weather models is the requirement for a tangent linear and adjoint version of the full nonlinear model. The adjoint is an extremely useful tool once developed (Errico, 1997); however, it requires many man-years to code and test, especially for complex models like ENLIL. For less complex MHD models, however, it may be possible to code an adjoint for use in solar wind forecasting. One of the main advantages of ensemble-based methods is their ease of implementation on a new problem.




Figure 5. Polar plots of the difference between the absolute error in the LETKF analysis and model-only ensemble. Lower absolute errors in the LETKF analysis are denoted by positive values (the red regions), and greater absolute errors in the LETKF analysis are denoted by negative values (the blue regions).




Figure 6. Hovmoller plots showing the difference between the absolute errors of the LETKF analysis and the pure ensemble for varying radial coordinates along the Earth-Sun line. Positive values (red) indicate that the LETKF has a lower absolute error, and negative values (blue) indicate that the model-only ensemble performs better.

Additionally, variational methods require that a background error covariance matrix is provided for each assimilation window, in other words, the uncertainty regarding the initial condition supplied for use in the data assimilation scheme. This matrix is generated by computing the errors of decades worth of forecasts and reanalysis data, something that is not available to us in the space weather field. For sequentially based methods like the ensemble Kalman filters, this matrix only has to be provided at the start of the data assimilation effort but is propagated by the ensemble from then on.

Similarly, as variational methods rely on linearizations, the data assimilation window cannot be too long to avoid the buildup of strong nonlinearities. In numerical weather prediction a typical window length is 6 to 12 h, and it is unclear how long the window can be for space weather applications. It will depend on the spatial resolution of the model, as higher resolution typically means stronger nonlinearities with shorter growth time scales and also on how close we are to the true solution, which is how wide the covariances are. Numerical weather prediction at the global scale is quite accurate due to the enormous amount of relatively accurate observations in each 6 h window, presently close to 10^7 observations. This means that the initial guess over the assimilation will be quite accurate too, so that although the full system is very nonlinear, we are so close to the true solution that linearizations work well. Space weather seems to be a long way away from this situation. (This does not mean that variational methods should be abandoned at this stage as one can run an ensemble of variational methods to explore the posterior pdf, similar to an ensemble smoother. This is a technique also used in weather prediction.)

Ensemble Kalman filters such as the one explored in this paper do rely on a Gaussian assumption of the prior pdf and on the likelihood. These will be violated, but, interestingly enough, ensemble Kalman filters have been found to be quite robust to non-Gaussian situations. Although not completely understood, this might be related to the fact that another way to derive the Kalman filter, and to some extent the ensemble Kalman filters, is to assume that the best estimate is a linear combination of the prior best estimate and the observations. That is often true to first order.

Ensemble Kalman filters need localization. The localization scheme used in these experiments was a relatively simple asymmetric distance-based localization on the $\mathbf{R}$ matrix with a localization radius of 0.01 AU. This localization radius may not be an accurate length scale, and different variables may require different localization radii (i.e., spurious correlations may need to be removed between radial momentum and longitudinal magnetic field, for example). Additionally, this distance-based localization scheme is perhaps not ideal for the solar wind simulations, as the variables are highly correlated along the magnetic field lines but weakly correlated perpendicular to them, due to the “freezing in” of the plasma and magnetic field. Therefore, an anisotropic localization scheme that follows the magnetic field lines in the domain may be more realistic and allow for better assimilation of the observations. Alternatively, if the magnetic field lines cannot be accurately ascertained from the model, then an anisotropic localization scheme may be more suitable that follows an Archimedean spiral radially outward from the Sun in the steady flow. When a CME is present the localization area should probably be closer to an expanding arc. More research is required to adequately analyze appropriate localization methods for the solar wind.

The magnetic field has the strong physical constraint, $\nabla \cdot \mathbf{B} = 0$, which needs to be conserved by the data assimilation analysis. For ENLIL, this is difficult to enforce at run time as the model solves for unipolar magnetic field, to minimize numerical diffusion effects. In ENLIL, $\nabla \cdot \mathbf{B} = 0$ is assumed at the previous time step and then evolved forward in time such that $\nabla \cdot \mathbf{B}$ remains 0 at the next time step. However, if $\nabla \cdot \mathbf{B} \neq 0$ at the previous time step, as a result of the data assimilation breaking the balance present in the system, then ENLIL will not enforce $\nabla \cdot \mathbf{B} = 0$. This will introduce numerical instabilities that may cause the model to become unstable and/or unphysical at future time steps. This has been verified by experiments performed in Lang (2016).

To get around this issue, the $\mathbf{B}$ field is removed from the state vector such that the magnetic field is not updated by the data assimilation scheme. This can be used even when using observations of the magnetic field because the LETKF will use the ensemble-based correlations between the magnetic field and the other state variables to update those. Information of the magnetic field is then passed on from the updated variables to the magnetic field itself through the ENLIL model evolution, in a manner similar to how CMEs are introduced into the ENLIL model. For a MHD code that enforces $\nabla \cdot \mathbf{B} = 0$ at each time step (e.g., by solving for the vector potential of the magnetic field), the data assimilation does not have to ensure that condition, but we can call the model routine to enforce it immediately after the data assimilation step.

Another area that needs attention is the large range of the variable values between a quiet phase and when a CME is present. When the observations suddenly indicate that a CME is present, a huge change to the model state is enforced by the data assimilation. This change can easily push the model out of balance, resulting in numerical instabilities and potential model instability. Similar problems are encountered in weather prediction, for instance when clouds are detected in observations, but the model does not have clouds at those grid points. A cloud represents a completely different thermodynamic state of the system, which is a large change in several model variables, and this is still an outstanding problem in numerical weather prediction. A potential solution is to use a smoother, so a method that updates the present state using observations from the future, allowing for a smoother model trajectory. A difficulty in space weather is that the sparsity of observations might preclude this solution.

In the experiments performed in Lang (2016), all ensemble members tended to collapse on one model evolution, and the LETKF became ineffective. This was due to the use of the same boundary conditions at the inner boundary, forcing the model evolutions to become similar quickly due to the highly driven nature of the solar wind. This was especially problematic when a CME was detected in the observations as none of the ensemble members contained a CME. This means that it is vital that different time-varying boundary conditions are used in order for ensemble-based methods to work. Again, much work is needed to gain experience in what the best way is to solve this problem. Each of these varying boundary conditions has to be a potential realistic realization of the real boundary condition, which suggests, for instance, that at least a few ensemble members should have a CME at the inner boundary at varying degrees of maturity.

This initial study has assumed that there is no model error present in the ENLIL model, which is obviously a poor assumption. These model errors typically fall under two main categories: systematic model errors and stochastic model error. Systematic model errors can be composed of errors in the model and biases that are constant over the assimilation window. Any known biases/systematic errors present in the prior model should be removed before the data assimilation is performed, and if they are not known but assumed to be present, they should be estimated. One possible method of diagnosing and removing functional model errors was

proposed by Lang et al. (2016), which use differences between a data assimilation run and a pure model run as a proxy for the model error present. Several methods have been proposed to estimate model biases, such as Dee (2005) and Tremolet (2006) among many others. To incorporate stochastic model uncertainty, it is necessary to generate a stochastic model error covariance matrix, $\mathbf{Q}$, that contains the relevant MHD balances to perturb the ensemble with a model error term at every time step and improve the quality of the assimilation. It is unclear what this matrix should look like, but it is important to realize that it should not contain the statistics of the missing physics per se but its influence on the resolved scales. One possible method of doing this would be to start from an initial guess and learn from the data assimilation experiments what the best formulation is, perhaps exploring techniques from machine learning. The first guess might come from a scaled down version, both in length scale and in amplitude, of a covariance matrix generated from samples of a long model run.

8. Conclusions

This is the first study to incorporate advanced data assimilation methods into solar wind models, building upon the work started in Lang (2016), and these experiments show that the LETKF is a useful tool that has potential for estimating the solar wind. It can be seen that the LETKF can improve upon the model-only ensemble at the observation point, and those improvements are advected radially outward. However, due to the solar wind being a highly driven system dominated by the dynamics of the Sun, the improvements made by the LETKF at the observation are confined to the magnetic flux tube on which the observation is located.

The density and radial momentum variables were very responsive at the observation point to the data assimilation analysis, moving close to the truth and returning to the background state as this improvement is advected away. However, the sharp gradient caused by this improvement is causing numerical wavelike patterns to form. Further research is required to determine the most effective solution.

The magnetic field vector, however, is not very responsive and does not appear to respond to the changes in the density, temperature, and momentum variables. While there are improvements in the magnetic field variables as a result of using the LETKF (especially in the latitudinal and longitudinal directions, not shown here), it is not possible to directly assimilate magnetic field observations while enforcing the $\nabla \cdot \mathbf{B} = 0$ criterion, either by one of the two methods suggested in Appendix B, or by rigorously enforcing $\nabla \cdot \mathbf{B} = 0$ at each time step in the numerical model.

While there are many challenges that still need to be overcome in solar wind data assimilation, these experiments show that data assimilation has great potential for the improvement of solar wind forecasting.

Appendix A: Local Ensemble Transform Kalman Filter

The Local Ensemble Transform Kalman Filter (LETKF) was developed by Hunt et al. (2007), based upon the Local Ensemble Kalman Filter (LEKF) developed by Ott et al. (2004). In the LETKF observation localization is performed. The principle behind the LETKF is to perform the analysis at the grid points in parallel, updating the state at each grid point using only observations contained within a local region of that grid point. The LETKF is an extension of the Ensemble Transform Kalman Filter (ETKF) (Bishop et al., 2001; Majumdar et al., 2002), which transforms the EnKF algorithm into the (typically) lower dimensional ensemble space, allowing more efficient computations.

The LETKF transforms the problem in state space to ensemble space as follows. At each observation time a perturbation matrix is generated for the state, which is defined as

$$\mathbf{X} = (\mathbf{x}^{(1)} - \bar{\mathbf{x}}, \dots, \mathbf{x}^{(M)} - \bar{\mathbf{x}}), \quad (\text{A1})$$

and in addition to this, an ensemble of model observations is generated, such that

$$\mathbf{Y} = (\mathbf{y}^{(1)} - \bar{\mathbf{y}}, \dots, \mathbf{y}^{(M)} - \bar{\mathbf{y}}), \quad (\text{A2})$$

where $\mathbf{y}^{(m)} = H(\mathbf{x}^{(m)})$ and $\bar{\mathbf{y}} = \frac{1}{M} \sum_{m=1}^M [H_k(\mathbf{x}^{(m)})]$.

In the LETKF, localization is performed on the observation error covariance matrix, $\mathbf{R}$, therefore limiting the observations that affect the state in the analysis. This is done by applying a “forgetting factor” (Nerger et al., 2012), v_r , to $\mathbf{R}$ for each observation, such that the observations further away from the analysis grid point get

less weight than those closer to the analysis grid point. For example, v could be defined by a distance-based Gaussian function between the analysis grid point and the observation grid point. The localized observation error covariance matrix shall be denoted by $\tilde{\mathbf{R}}$, where

$$\tilde{\mathbf{R}} = \frac{1}{v} \mathbf{R}. \quad (\text{A3})$$

A matrix $\mathbf{C} \in \mathbb{R}^{M \times N_y}$ is defined such that

$$\mathbf{C} = \mathbf{Y}^T \tilde{\mathbf{R}}^{-1}, \quad (\text{A4})$$

which is used to calculate the matrix, $\mathbf{T} \in \mathbb{R}^{M \times M}$:

$$\mathbf{T} = \left(\frac{M-1}{1+\rho} \mathbf{I} + \mathbf{C}\mathbf{Y} \right)^{-1}, \quad (\text{A5})$$

where $\rho > 0$ is the multiplicative covariance inflation factor. Typical values range from 0.01 up to 10 in extreme cases.

The weighting matrix $\mathbf{W} \in \mathbb{R}^{M \times M}$ and mean weighting vector $\bar{\mathbf{w}} \in \mathbb{R}^M$ are computed as

$$\bar{\mathbf{w}} = \mathbf{T}\mathbf{C}(\mathbf{y} - \bar{\mathbf{y}}) \quad (\text{A6})$$

$$\mathbf{W} = ((M-1)\mathbf{T})^{\frac{1}{2}}. \quad (\text{A7})$$

An ensemble of weight vectors $\mathbf{W}^{\mathbf{a}(M)}$ is created by adding $\bar{\mathbf{w}}^a$ to the columns of $\mathbf{W}^a$, such that the weight vector for the m th ensemble member is given by

$$\mathbf{w}^{\mathbf{a}(m)} = \bar{\mathbf{w}} + \mathbf{W}_m, \quad (\text{A8})$$

where $\mathbf{W}_m$ is the m th column of $\mathbf{W}$.

Then the m th analysis ensemble member is given by

$$\mathbf{x}^{\mathbf{a}(m)} = \bar{\mathbf{x}} + \mathbf{X}\mathbf{w}^{\mathbf{a}(m)}. \quad (\text{A9})$$

Appendix B: Ensuring $\nabla \cdot \mathbf{B} = 0$

In this appendix several methods are presented and discussed that enforce $\nabla \cdot \mathbf{B} = 0$ in each updated ensemble member.

1. Enforce $\nabla \cdot \mathbf{B} = 0$ after the LETKF analysis via a projection method.

Via the Helmholtz decomposition (Arfken et al., 2011), any field can be written as the sum of the curl of another field $\mathbf{A}$ and the gradient of a scalar field, φ . Therefore, $\mathbf{B}$ can be written as

$$\mathbf{B} = \nabla \times \mathbf{A} + \nabla \varphi. \quad (\text{B1})$$

Taking the divergence of equation (B1), it can be seen that

$$\nabla \cdot \mathbf{B} = \nabla \cdot \nabla \varphi = \nabla^2 \varphi. \quad (\text{B2})$$

This is nonzero if $\nabla \varphi$ is nonzero. Therefore, in order to preserve the $\nabla \cdot \mathbf{B}$ condition, $\nabla \varphi$, must be subtracted from the DA analysis of the magnetic field variables.

To do this, it is necessary to calculate $\nabla \cdot \mathbf{B}$ and then solve the Poisson equation

$$\nabla^2 \varphi = \nabla \cdot \mathbf{B} \quad (\text{B3})$$

for φ .

The issue with this is that the Poisson equation is an elliptic equation that typically has to be solved iteratively, which can be expensive to compute for high-dimensional systems.

2. Encode $\nabla \cdot \mathbf{B} = 0$ into the Kalman filter update equations via the covariances.

For the second method, we start by looking at the ETKF update equation for the magnetic field component of the state vector, which is given by

$$\mathbf{B}_i^{\mathbf{a}(\mathbf{M})} = \mathbf{B}_i^{\mathbf{f}(\mathbf{M})} + \sum_{m=1}^N \left[\left(\mathbf{B}_i^{\mathbf{f}(\mathbf{m})} - \overline{\mathbf{B}_i^{\mathbf{f}(\mathbf{M})}} \right) W_i^{B(m)} \right], \quad (\text{B4})$$

where $\mathbf{B}_i^{\mathbf{f}(\mathbf{M})}$ and $\mathbf{B}_i^{\mathbf{a}(\mathbf{M})}$ are the forecast and analysis magnetic field ensemble, each with M ensemble members, at time step i , and $W_i^{B(m)}$ are the weights generated for the m th ensemble member of the magnetic field variable, by the ETKF methodology.

By taking the divergence of the analysis state, we obtain the following equation:

$$\nabla \cdot \mathbf{B}_i^{\mathbf{a}(\mathbf{M})} = \nabla \cdot \mathbf{B}_i^{\mathbf{f}(\mathbf{M})} + \sum_{m=1}^M \left[\left(\nabla \cdot \mathbf{B}_i^{\mathbf{f}(\mathbf{m})} - \nabla \cdot \overline{\mathbf{B}_i^{\mathbf{f}(\mathbf{M})}} \right) W_i^{B(m)} \right] = 0. \quad (\text{B5})$$

Using $\nabla \cdot \mathbf{B}_i^{\mathbf{f}(\mathbf{m})} = 0$ for each ensemble member from the prior model, and the prior mean is also divergence free because the mean is a linear operator, implying $\nabla \cdot \overline{\mathbf{B}_i^{\mathbf{f}(\mathbf{M})}} = 0$; it can be seen that $\nabla \cdot \mathbf{B}_i^{\mathbf{a}(\mathbf{M})} = 0$ for the ETKF. Therefore, the ETKF satisfies the balance equation, as required.

Unfortunately, when we localize the ETKF for the LETKF, $W_i^{B(m)}$ will become dependent upon the space domain, and hence the divergence of a term within the sum of equation (B4) becomes

$$\begin{aligned} \nabla \cdot \left(\left(\mathbf{B}_i^{\mathbf{f}(\mathbf{m})} - \overline{\mathbf{B}_i^{\mathbf{f}(\mathbf{M})}} \right) W_i^{B(m)} \right) &= \left(\nabla \cdot \mathbf{B}_i^{\mathbf{f}(\mathbf{m})} - \nabla \cdot \overline{\mathbf{B}_i^{\mathbf{f}(\mathbf{M})}} \right) W_i^{B(m)} \\ &\quad + \left(\mathbf{B}_i^{\mathbf{f}(\mathbf{m})} - \overline{\mathbf{B}_i^{\mathbf{f}(\mathbf{M})}} \right) \nabla \cdot W_i^{B(m)}, \end{aligned} \quad (\text{B6})$$

which implies

$$\nabla \cdot \mathbf{B}_i^{\mathbf{a}(\mathbf{M})} = \sum_{m=1}^M \left(\mathbf{B}_i^{\mathbf{f}(\mathbf{m})} - \overline{\mathbf{B}_i^{\mathbf{f}(\mathbf{M})}} \right) \nabla \cdot W_i^{B(m)}. \quad (\text{B7})$$

To ensure that this equation is equal to 0, it is required that $\nabla \cdot W_i^{B(m)} = 0$ for all ensemble members. This is of similar complexity to solving $\nabla \cdot \mathbf{B}_i^{\mathbf{a}(\mathbf{m})} = 0$ directly for each ensemble member, so it appears that little has been gained from this.

3. Solve a variational problem for each ensemble member.

Another potential solution is to realize that the ETKF update equation for each ensemble member can be found as the solution to a variational problem. The $\nabla \cdot \mathbf{B}_i^{\mathbf{a}(\mathbf{m})} = 0$ condition can be enforced via a Lagrange multiplier in that variational problem. We will not elaborate on that solution here.

Acknowledgments

The data used are listed in the references, tables, and figure captions. The boundary condition data used to initialize the runs are generated from WSA model runs and were provided to us by the Met Office; however, these data can also be freely downloaded from the repositories at Space Weather Prediction Center (SWPC) (2017). This work was funded by the Natural Environment Research Council under NE/J005878/1. In addition, P.J.v.L. thanks the European Research Council (ERC) for funding the CUNDA project under the European Union's Horizon 2020 research and innovation program. This work used the ARCHER UK National Supercomputing Service (<http://www.archer.ac.uk>).

References

- Amezcu, J., Goodliff, M., & van Leeuwen, P. J. (2017). A weak constraint 4D EnsembleVar. Part I: Formulation and simple model experiments. *Tellus A*, 69(1), 1271564.
- Anderson, J. (2007). Exploring the need for localization in ensemble data assimilation using a hierarchical ensemble filter. *Physica D: Nonlinear Phenomena*, 230(1), 99–111.
- Anderson, J., Hoar, T., Raeder, K., Liu, H., Collins, N., Torn, R., & Avellano, A. (2009). The data assimilation research testbed: A community facility. *Bulletin of the American Meteorological Society*, 90(9), 1283.
- Arfken, G., Weber, H., & Harris, F. (2011). *Mathematical Methods for Physicists: A Comprehensive Guide*. London: Academic press.
- Arge, C., & Pizzo, V. (2000). Improvement in the prediction of solar wind conditions using near-real time solar magnetic field updates. *Journal of Geophysical Research*, 105(A5), 10,465–10,479.
- Bishop, C., Etherton, B., & Majumdar, S. (2001). Adaptive sampling with the ensemble transform Kalman filter: Part I: Theoretical aspects. *Monthly Weather Review*, 129(3), 420–436.
- Board, S. S. (2008). *Severe Space Weather Events—Understanding Societal and Economic Impacts: A Workshop Report*. Washington, DC: National Academies Press.
- Browne, P., & Wilson, S. (2015). A simple method for integrating a complex model into an ensemble data assimilation system using MPI. *Environmental Modelling & Software*, 68, 122–128.
- Burt, J., & Smith, B. (2012). Deep space climate observatory: The DSCOVR mission. In *Aerospace Conference, 2012 IEEE* (pp. 1–13). Big Sky, MT: IEEE.
- Cannon, P., Angling, M., Barclay, L., Curry, C., Dyer, C., Edwards, R., ... Underwood, C. (2013). *Extreme Space Weather: Impacts on Engineered Systems and Infrastructure*. London: Royal Academy of Engineering.
- Cranmer, S. (2008). On competing models of coronal heating and solar wind acceleration: The debate in '08. arXiv:0804.3058.
- Davis, C., Davies, J., Lockwood, M., Rouillard, A., Eyles, C., & Harrison, R. (2009). Stereoscopic imaging of an Earth-impacting solar coronal mass ejection: A major milestone for the Stereo mission. *Geophysical Research Letters*, 36, L08102. <https://doi.org/10.1029/2009GL038021>

- Dee, D. P. (2005). Bias and data assimilation. *Quarterly Journal of the Royal Meteorological Society*, 131(613), 3323–3343.
- De Moortel, I., & Browning, P. (2015). Recent advances in coronal heating. Retrieved from <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4410557/>
- Errico, R. (1997). What is an adjoint model? *Bulletin of the American Meteorological Society*, 78(11), 2577–2591.
- Evensen, G., Dee, D., & Schröter, J. (1998). Parameter estimation in dynamical models. In *Ocean Modeling and Parameterization* (pp. 373–398). Dordrecht, Netherlands: Springer.
- Eyles, C., Harrison, R., Davis, C., Waltham, N., Shaughnessy, B., Mapson-Menard, H., ... Rochus, P. (2009). The heliospheric imagers onboard the Stereo mission. *Solar Physics*, 254(2), 387–445.
- Fairbairn, D., Pring, S., Lorenc, A., & Roulstone, I. (2014). A comparison of 4DVar with ensemble data assimilation methods. *Quarterly Journal of the Royal Meteorological Society*, 140, 281–294.
- Gloeckler, G., Balsiger, H., Bürgi, A., Bochsler, P., Fisk, L., Galvin, A., ... Holzer, T. (1995). The solar wind and suprathermal ion composition investigation on the Wind spacecraft. *Space Science Reviews*, 71(1–4), 79–124.
- Goodliff, M., Amezcua, J., & van Leeuwen, P. J. (2017). A weak constraint 4DensembVar. Part II: Experiments with larger models. *Tellus A*, 69, 1271565.
- Gosling, J. T. (1993). The solar flare myth. *Journal of Geophysical Research*, 98(A11), 18,937–18,949. <https://doi.org/10.1029/93JA01896>
- Gropp, W., Lusk, E., Doss, N., & Skjellum, A. (1996). A high-performance, portable implementation of the MPI message passing interface standard. *Parallel computing*, 22(6), 789–828.
- Hassler, D., Dammasch, I., Lemaire, P., Brekke, P., Curdt, W., Mason, H., ... Wilhelm, K. (1999). Solar wind outflow and the chromospheric magnetic network. *Science*, 283(5403), 810–813.
- Hodyss, D., & Nichols, N. K. (2015). The error of representation: Basic understanding. *Tellus A*, 67, 24822.
- Houtekamer, P., & Mitchell, H. (2001). A sequential ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 129(1), 123–137.
- Hunt, B., Kostelich, E., & Szunyogh, I. (2007). Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter. *Physica D: Nonlinear Phenomena*, 230(1), 112–126.
- Jackson, B. (1985a). Helios observations of the earthward-directed mass ejection of 27 November, 1979. *Solar physics*, 95(2), 363–370.
- Jackson, B. (1985b). Imaging of coronal mass ejections by the Helios spacecraft. *Solar physics*, 100(1–2), 563–574.
- Kaiser, M., Kucera, T., Davila, J., Cyr, O. S., Guhathakurta, M., & Christian, E. (2008). The Stereo mission: An introduction. *Space Science Reviews*, 136, 5–16.
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering*, 82, 35–45.
- Kalman, R., & Bucy, R. (1961). New results in linear filtering and prediction theory. *Journal of Fluids Engineering*, 83(1), 95–108.
- Kivelson, M., & Russell, C. (1995). *Introduction to space physics*. Cambridge, UK: Cambridge University Press.
- Lang, M. (2016). Model improvement using data assimilation (PhD Thesis), University of Reading.
- Lang, M., van Leeuwen, P. J., & Browne, P. A. (2016). A systematic method of parameterisation estimation using data assimilation. *Tellus A*, 68, 29012.
- Leslie, J., & Cole, S. (2016). DSCOVR: Deep Space Climate Observatory. Retrieved from <https://www.nesdis.noaa.gov/content/dscovr-deep-space-climate-observatory>
- Lopez, R. (1987). Solar cycle invariance in solar wind proton temperature relationships. *Journal of Geophysical Research*, 92(A10), 11,189–11,194.
- Majumdar, S., Bishop, C., Etherton, B., & Toth, Z. (2002). Adaptive sampling with the ensemble transform Kalman filter: Part II: Field program implementation. *Monthly Weather Review*, 130(5), 1356–1369.
- Manoharan, P., & Ananthakrishnan, S. (1990). Determination of solar-wind velocities using single-station measurements of interplanetary scintillation. *Monthly Notices of the Royal Astronomical Society*, 244, 691–695.
- McComas, D. J., Velli, M., Lewis, W. S., Acton, L. W., Balat-Pichelin, M., Bothmer, V., ... Zurbuchen, T. H. (2007). Understanding coronal heating and solar wind acceleration: Case for in situ near-Sun measurements. *Reviews of Geophysics*, 45, RG1004. <https://doi.org/10.1029/2006RG000195>
- Miyoshi, T., Sato, Y., & Kadowaki, T. (2010). Ensemble Kalman filter and 4D-Var intercomparison with the Japanese operational global analysis and prediction system. *Monthly Weather Review*, 138(7), 2846–2866.
- Nakamura, G., & Potthast, R. (2015). *Inverse modeling* (pp. 2053–2563). Bristol: IOP Publishing. <https://doi.org/10.1088/978-0-7503-1218-9>
- Nerger, L., Hiller, W., & Schröter, J. (2005). PDAF—The parallel data assimilation framework: Experiences with Kalman filtering, *Use of high performance computing in meteorology: Proceedings of the Eleventh ECMWF Workshop on the Use of High Performance Computing in Meteorology, Reading, UK, 25–29 October 2004/Eds.: Walter Zwielfhofer; Geoge Mozdzyński, Singapore [ua]: World* (pp. 63–83). London: World Scientific.
- Nerger, L., Janjić, T., Schröter, J., & Hiller, W. (2012). A regulated localization scheme for ensemble-based Kalman filters. *Quarterly Journal of the Royal Meteorological Society*, 138(664), 802–812.
- Odstrcil, D. (2003). Modeling 3-D solar wind structure. *Advances in Space Research*, 32(4), 497–506.
- Odstrcil, D. (2004). Enlil: A numerical code for solar wind disturbances. Retrieved from http://ccmc.gsfc.nasa.gov/models/Code_description.pdf
- Odstrcil, D., & Pizzo, V. (1999). Distortion of the interplanetary magnetic field by three-dimensional propagation of coronal mass ejections in a structured solar wind. *Journal of Geophysical Research*, 104(A12), 28,225–28,239.
- Odstrcil, D., Riley, P., & Zhao, X. (2004). Numerical simulation of the 12 May 1997 interplanetary CME event. *Journal of Geophysical Research*, 109, A02116. <https://doi.org/10.1029/2003JA010135>
- Ott, E., Hunt, B., Szunyogh, I., Zimin, A., Kostelich, E., Corazza, M., ... Yorke, J. (2004). A local ensemble Kalman filter for atmospheric data assimilation. *Tellus A*, 56(5), 415–428.
- Owens, M., & Forsyth, R. (2013). The heliospheric magnetic field. *Living Reviews in Solar Physics*, 10, 5.
- Owens, M., Horbury, T., Wicks, R., McGregor, S., Savani, N., & Xiong, M. (2014). Ensemble downscaling in coupled solar wind-magnetosphere modeling for space weather forecasting. *Space Weather*, 12, 395–405. <https://doi.org/10.1002/2014SW001064>
- Parker, E. (1958). Dynamics of the interplanetary gas and magnetic fields. *The Astrophysical Journal*, 128, 664–676.
- Parsons, A., Biesecker, D., Odstrcil, D., Millward, G., Hill, S., & Pizzo, V. (2011). Wang-Sheeley-Argé–Enlil cone model transitions to operations. *Space Weather*, 9, 3004. <https://doi.org/10.1029/2011SW000663>
- Reich, S., & Cotter, C. (2015). *Probabilistic forecasting and Bayesian data assimilation*. Cambridge: Cambridge University Press.
- Riley, P., Linker, J., & Arge, C. (2015). On the role played by magnetic expansion factor in the prediction of solar wind speed. *Space Weather*, 13, 154–169. <https://doi.org/10.1002/2014SW001144>

- Smith, P., Dance, S., Baines, M., Nichols, N., & Scott, T. (2009). Variational data assimilation for parameter estimation: Application to a simple morphodynamic model. *Ocean Dynamics*, 59(5), 697–708.
- Space Weather Prediction Center (SWPC) (2017). Wang-Sheeley-Arge (WSA)-Enlil solar wind prediction, NOAA National Centers for Environmental Information. Retrieved from <https://www.ngdc.noaa.gov/enlil/>, <https://doi.org/10.7289/V5445JGH>
- Stone, E., Frandsen, A., Mewaldt, R., Christian, E., Margolies, D., Ormes, J., & Snow, F. (1998). The Advanced Composition Explorer: The Advanced Composition Explorer mission. *Space Science Reviews*, 86, 1–22.
- Szunyogh, I., Kostelich, E., Gyarmati, G., Kalnay, E., Hunt, B., Ott, E., . . . Yorke, J. (2008). A local ensemble transform Kalman filter data assimilation system for the NCEP global model. *Tellus A*, 60(1), 113–130.
- Tremolet, Y. (2006). Accounting for an imperfect model in 4D-Var. *Quarterly Journal of the Royal Meteorological Society*, 132(621), 2483–2504.
- van Leeuwen, P. J. (2015). Representation errors and retrievals in linear and nonlinear data assimilation. *Quarterly Journal of the Royal Meteorological Society*, 141, 1612–1623. <https://doi.org/10.1002/qj.2464>
- van Leeuwen, P. J., Cheng, Y., & Reich, S. (2015). *Nonlinear data assimilation*. Switzerland: Springer. <https://doi.org/10.1007/978-3-319-18347-3>
- Wang, Y., & Sheeley Jr., N. (1992). On potential field models of the solar corona. *The Astrophysical Journal*, 392, 310–319.
- Xie, H., Ofman, L., & Lawrence, G. (2004). Cone model for halo CMEs: Application to space weather forecasting. *Journal of Geophysical Research*, 109, A03109. <https://doi.org/10.1029/2003JA010226>