AGU PUBLICATIONS

Space Weather

RESEARCH ARTICLE
10.1002/2017SW001752

Special Section:
Low Earth Orbit Satellite Drag: Science and Operational Impact

An Ensemble Kalman Filter for the Thermosphere-Ionosphere

S. M. Codrescu[1], M. V. Codrescu[2], and M. Fedrizzi[1,2]

[1]Cooperative Institute for Research in Environmental Science, University of Colorado Boulder, Boulder, CO, USA, [2]Space Weather Prediction Center, NOAA, Boulder, CO, USA

Key Points:
- Improved Thermospheric neutral density specification using an Ensemble Kalman Filter for quiet conditions
- An ensemble distributed via randomized forcing produces sufficient covariance for assimilation in quiet conditions
- Covariance localization unnecessary for an ensemble distributed via randomized system forcing

Correspondence to:
S. M. Codrescu,
stefan.codrescu@colorado.edu

Citation:
Codrescu, S. M., Codrescu, M. V., & Fedrizzi, M. (2018). An Ensemble Kalman Filter for the thermosphere-ionosphere. *Space Weather*, *16*, 57–68. https://doi.org/10.1002/2017SW001752

Received 9 AUG 2017
Accepted 30 DEC 2017
Accepted article online 8 JAN 2018
Published online 18 JAN 2018

Abstract A principal limitation of physics-based modeling in the Thermosphere-Ionosphere system is the large uncertainty associated with the implementation of the external forcing of the system, including high-latitude convection and particle precipitation, solar UV/EUV fluxes, and waves propagating from below. As measuring these quantities with sufficient spatial and temporal resolution is prohibitively costly, a more practical approach to improve model results is to assimilate more readily available measurements of the system. We discuss considerations in implementing an Ensemble Kalman Filter (EnKF) for a strongly forced system versus a chaotic system, overview an EnKF implementation for the strongly forced Coupled Thermosphere, Ionosphere, Plasmasphere, and Electrodynamics model, and present encouraging improvements to neutral density specification obtained by assimilating Challenging Minisatellite Payload (CHAMP) neutral density measurements. The model results show improvement in comparisons with both CHAMP and Gravity Recovery and Climate Experiment (GRACE) measurements during a geomagnetically quiet period at solar minimum.

1. Introduction

Neutral density is of particular importance to orbit determination for satellites in Low Earth Orbit (LEO). Atmospheric drag is a function of neutral density; consequently, a large contribution to the uncertainty in tracking space debris, forecasting satellite collisions, and projecting satellite lifetimes can be attributed to uncertainty in specification of the global neutral density field. Changing space weather conditions are the largest contributor to the changes in the global neutral density field. Physics-based models of the Thermosphere-Ionosphere region encompassing LEO are the most promising for their ability to propagate state in time using system dynamics (Fedrizzi et al., 2012; Moe & Moe, 2011). State propagation is critical when measurements are sparse—the measurements available at a given point in time can be used not only to correct the immediate state but may also improve future states by guiding the trajectory of the model in state space. Empirical models are limited in their ability to represent structure and, in practice, suffer during disturbed periods when the Thermosphere-Ionosphere exhibits large deviations from climatology.

The ability of physics-based numerical models of the Thermosphere-Ionosphere to reproduce observations and to make accurate specification of the current and future state of the system is at present limited by incomplete knowledge of the system forcing as a function of time (Codrescu et al., 1997, 1989; Crowley et al., 2006; Fuller-Rowell et al., 2000). The use of a proxy index for solar Ultra Violet (UV) and Extreme Ultra Violet (EUV) radiation, the use of statistical patterns of particle precipitation and plasma convection for the calculation of Joule heating, and the use of an hourly averaged parameterization of the lower boundary condition lead to large uncertainties in model predictions, especially during disturbed conditions. The cost of measuring the necessary system forcing parameters, notably Joule heating, with sufficient spatial and temporal resolution, is prohibitive. A viable alternative is to use a Data Assimilation (DA) scheme to correct the state of a physics-based numerical model according to more easily obtained measurements containing sufficient information to infer system forcing.

Currently, we develop and extend an Ensemble Kalman Filter (EnKF) architecture first explored in Codrescu et al. (2004). The EnKF is a computationally feasible Monte Carlo approximation to the Kalman Filter, useful for large, nonlinear systems (Evensen, 1994). For the Thermosphere-Ionosphere, the state covariance, that is, the cross correlation between state elements, is nonstationary. The EnKF is a practical choice since the necessary covariance can be determined at each assimilation time step from statistics across an ensemble of model simulations (members). While the EnKF has found wide use in both terrestrial and oceanographic applications,

©2018. American Geophysical Union. All Rights Reserved.

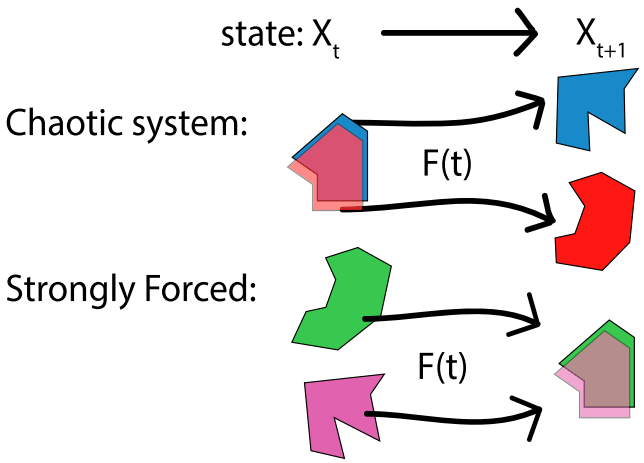

Figure 1. An illustration of difference between a chaotic versus a strongly forced system. Strongly forced systems are most sensitive to the forcing over time, while chaotic systems are more sensitive to the initial condition.

unfortunately the sparsity of measurements and the nature of the strongly forced systems present in space science have delayed the similar wide application of data assimilation in research and operational Space Weather. The results described in this study are a step in the direction of using numerical models for specification and, in the future, forecast of total mass density in an operational environment to help tracking of space debris and prediction of satellite orbits.

In this work, we introduce an implementation of an Ensemble Kalman Filter for the Thermosphere-Ionosphere. First, we motivate our approach by examining the difference between strongly forced and chaotic systems. In sections 3 and 4, we present the specifics of an implementation utilizing the Coupled Thermosphere, Ionosphere, Plasmasphere, and Electrodynamics (CTIPe) model as the forecast model. In section 5, we present an experiment assimilating CHAMP satellite observations of neutral density and comparison of the result with measurements of neutral density from the GRACE satellite. Finally, we close with a discussion of results in the context of recent DA work in the Thermosphere-Ionosphere and future research directions.

2. System Types

The goal of the ensemble within the EnKF is to maintain a distribution of points in state space representative of the uncertainty in the model state with respect to reality. Concretely, a wider distribution of ensemble members indicates more uncertainty. The initial Evensen (1994) formulation of the EnKF, and following widespread adoption, initializes the ensemble distribution by randomizing the initial condition of each ensemble member. In other words, each member is seeded with an initial condition perturbed by adding a random component sampled from some distribution with relevant covariance. This is natural for chaotic systems as the principle uncertainty in the current state of the system is due to uncertainty in the initial condition.

Consider, however, the difference between a chaotic and a strongly forced system. Chaotic systems are characterized by a strong dependence on the initial condition. In a chaotic system, two slightly different initial conditions will diverge over time to widely different locations in state space. On the other hand, a strongly forced system exhibits a strong dependence on the forcing over time. Two significantly different initial conditions will converge toward a more similar state if each is given the same forcing over time (Figure 1).

Recognizing that the Thermosphere-Ionosphere is strongly forced (Huba et al., 2014; Johnson & Killeen, 1995; KORENKOV, 2013), combined with the uncertainty of the available forcing inputs, we choose to create an ensemble by randomizing only the forcing applied to each ensemble member. The randomized forcing in time results in a distribution of ensemble members in state space that is conceptually similar to one created by randomizing the initial condition and using perturbed observations, yet better represents the uncertainty in modeling the Thermosphere-Ionosphere.

3. Ensemble Kalman Filter

The EnKF uses the familiar update equation from the Kalman filter equation (1) (Kalman, 1960) but approximates the Kalman gain K using $\hat{K}$, derived from an ensemble of state realizations, as will be shown below in equation (4).

$$x^a = x^f + K(y - h(x^f)) \tag{1}$$

Notationally, x is the EnKF state vector. Specifically, $x^f \in \mathbb{R}^n$ is the forecast state estimate. $x^a \in \mathbb{R}^n$ is the analysis estimate of the state taking into account measurements. K is the Kalman gain, $y \in \mathbb{R}^m$ is a vector of observations of the system, and $h(x) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a function mapping from state to measurements such that $y = h(x)$.

$\hat{K}$ is calculated from an ensemble containing q members, each an independently perturbed state realization where x^{f_i} is the EnKF state vector of the ith member of the ensemble and, similarly, $y^{f_i} = h(x^{f_i})$ are simulated observations according to the ith ensemble member.

Since the true state is unknown, the ensemble means $\bar{x}^f$ and $\bar{y}^f$ are used to approximate true state and measurements, respectively. The state error E_x and the observation error E_y are given by

$$E_x = [x^{f_1} - \bar{x}^f \dots x^{f_q} - \bar{x}^f] \text{ and } E_y = [y^{f_1} - \bar{y}^f \dots y^{f_q} - \bar{y}^f] \quad (2)$$

Then

$$P_{xy} = \frac{1}{q-1} E_x (E_y)^T \text{ and } P_{yy} = \frac{1}{q-1} E_y (E_y)^T \quad (3)$$

where P_{xy} is the state measurement error covariance and P_{yy} is the measurement error covariance. Finally,

$$\hat{K} = P_{xy}(P_{yy} + R)^{-1} \quad (4)$$

Once equation (1) has been evaluated at time t to obtain the analysis estimate x_t^a, the state can then be forecast into the future using the forecast model, creating x_{t+1}^f and the assimilation process repeats. For more depth on these equations, see Gillijns et al. (2006), noting the addition of observation error covariance $R \in \mathbb{R}^{m \times m}$ in equation (4).

4. Methodology

4.1. Forecast Model

The forecast model used in the assimilation scheme is the Coupled Thermosphere, Ionosphere, Plasmasphere, and Electrodynamics (CTIPe) model (Millward et al., 1996). CTIPe is a nonlinear physics-based numerical code. The model consists of four distinct components that are fully coupled. Included are a global thermosphere, a high-latitude ionosphere, a middle- and low-latitude ionosphere/plasmasphere, and an electrodynamics calculation of the global dynamo electric field. The latest developments include validation and tuning to reproduce observations (Codrescu et al., 2008, 2012; Fedrizzi et al., 2012).

The thermosphere code simulates the time-dependent structure of the wind vector, temperature, and density of the neutral thermosphere by numerically solving the nonlinear primitive equations of momentum, energy, and continuity (Fuller-Rowell & Rees, 1980). The global atmosphere is divided into a series of elements in geographic latitude, longitude, and pressure. Each grid point rotates with Earth to define a noninertial frame of reference in a spherical polar coordinate system. Latitude resolution is 2°, longitude resolution is 18°, and the vertical resolution is 1 scale height. In the vertical direction the atmosphere is divided into 15 levels, in logarithm of pressure, from a lower boundary of 1 Pa at 80 km altitude, to a varying altitude generally above 400 km.

In addition to the thermosphere, CTIPe includes fully dynamic and coupled descriptions of the high-latitude ionosphere and middle- and low-latitude ionosphere and plasmasphere. In these regions the model solves the relevant equations of continuity, momentum, and energy for O+ and H+ ions and electrons. The equations are solved along a large number of independent, magnetic flux tubes, orientated along an assumed eccentric dipole magnetic field.

The model inputs include a weighted average of $F_{10.7}$ nm radio flux, used for the calculation of heating, dissociation, and ionization rates though correlation with the Hinteregger reference spectrum (Hinteregger et al., 1981). Additionally, solar wind parameters are input at each time step: solar wind speed $|V_{sw}|$, solar wind density ρ_{sw}, magnitude of the solar wind interplanetary magnetic field in the GSE-YZ plane B_N, and solar wind magnetic field angle from $+Z$ in the GSE-YZ plane B_θ. Solar wind inputs determine the magnetospheric input to the model from empirical models of auroral precipitation (Fuller-Rowell & Evans, 1987) and convection electric fields (Weimer, 2005). The use of statistical models in the calculation of Joule heating introduces large uncertainties in the forcing of the model and consequently in the model results.

The EnKF state vector includes at each time step a subset of the system forcing that does not include the parameterization of tides, gravity waves, and planetary waves propagating into the Thermosphere from below. These are included in CTIPe by applying an hourly reproducible lower boundary condition determined from a climatological run of the Whole Atmosphere Model (WAM) (Akmaev, 2011; Fuller-Rowell et al., 2008). The influence of the lower boundary condition on CTIPe model results, although considerable, is expected to change on time scales of several days, not hours, and for this reason is not estimated in this study.

4.2. Assimilation Scheme

The assimilation scheme is implemented in C++ and uses the Eigen linear algebra library (Guennebaud et al., 2010) as well as the NetCDF file format to interface with the CTIPe model. The equations implemented

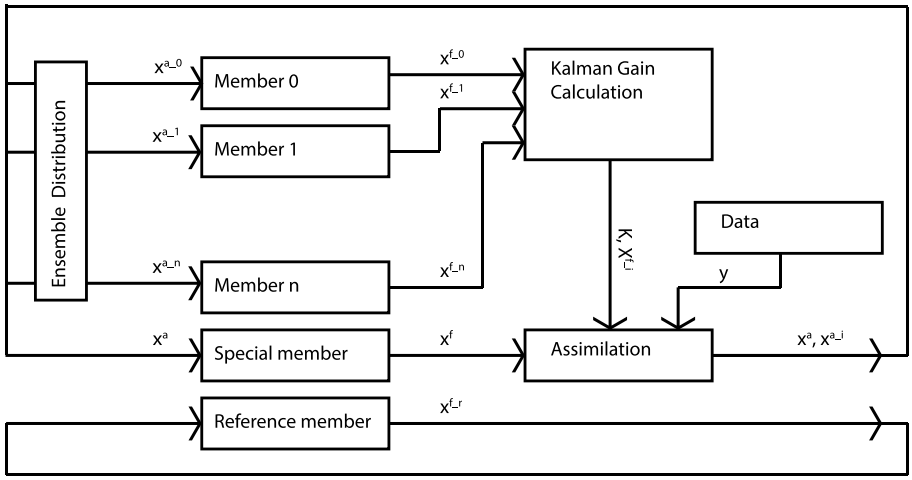

Figure 2. Visual layout of the major components and flow of information through our Ensemble Kalman Filter Data Assimilation Scheme.

are described in section 3. Importantly, this method calculates only the necessary terms of the covariance matrix. The code is primarily a research tool and, as such, is designed to be configurable, modular, and extensible. Figure 2 is a block diagram overviewing the configuration in which our results were obtained.

The ensemble is distributed by applying an independent and identically distributed (i.i.d.) set of input forcing parameters to each ensemble member. The i.i.d. input realizations are sampled from a normal distribution centered about a "special member." In other words, there exists an ensemble member whose forcing is equal to the mean of the distribution from which ensemble member forcings are sampled. Recall from equation (2) that the ensemble mean approximates the true state, and thus, the special member is a realization of the best estimate of the state at any time.

In addition, a reference member is run in parallel with no assimilation. It uses the initial condition and inputs that would have been available in an operational real time run (Codrescu et al., 2012) and is a useful basis to gauge the improvement due to DA.

5. Data Assimilation Experiment

5.1. Configurations

The assimilation scheme is run for the full day of 20 March 2007 – a geomagnetically quiet day near solar minimum in three otherwise identical configurations: (i) correction of the model state, (ii) model state correction and forcing estimation, and (iii) only forcing estimation. In all three, a 75 member ensemble with a 10 min assimilation time step is used. In each configuration, the ensemble is distributed at each assimilation time step by applying to each member's input a perturbation sampled from the following normal distributions: $F_{10.7} \sim \mathcal{N}(0, 50)$, $|V_{sw}| \sim \mathcal{N}(0, 50)$, $\rho_{sw} \sim \mathcal{N}(0, 5)$, $B_N \sim \mathcal{N}(0, 20)$, and $B_\theta \sim \mathcal{N}(0, 50)$.

The EnKF state vector is illustrated in equation (5) and consists of the model forcing parameters appended to model state. For the current experiment, the variables included in the state are those necessary to calculate neutral density, namely, neutral temperature, mass mixing ratios for atomic oxygen, molecular oxygen, and molecular nitrogen, mean molecular mass, as well as vertical and horizontal components of the neutral wind.

$$x = \begin{bmatrix} \text{model forcing} \\ \text{model state} \end{bmatrix} \tag{5}$$

Configuration (i) is a typical application of the EnKF in which the model state is estimated according to the measurements. Specifically, the model state component of the analysis EnKF state vector x^a is forecast by the background model at each assimilation time step. The model forcing components of x^a are ignored in this configuration. The forcing parameters used are the average over the assimilation time step for the special member, and a perturbation of that for each ensemble member.

In Configuration (ii), both the model state and the model forcing parameters are estimated. Similar to Configuration (i), the model state is forecast at each assimilation time step; however, the estimated model

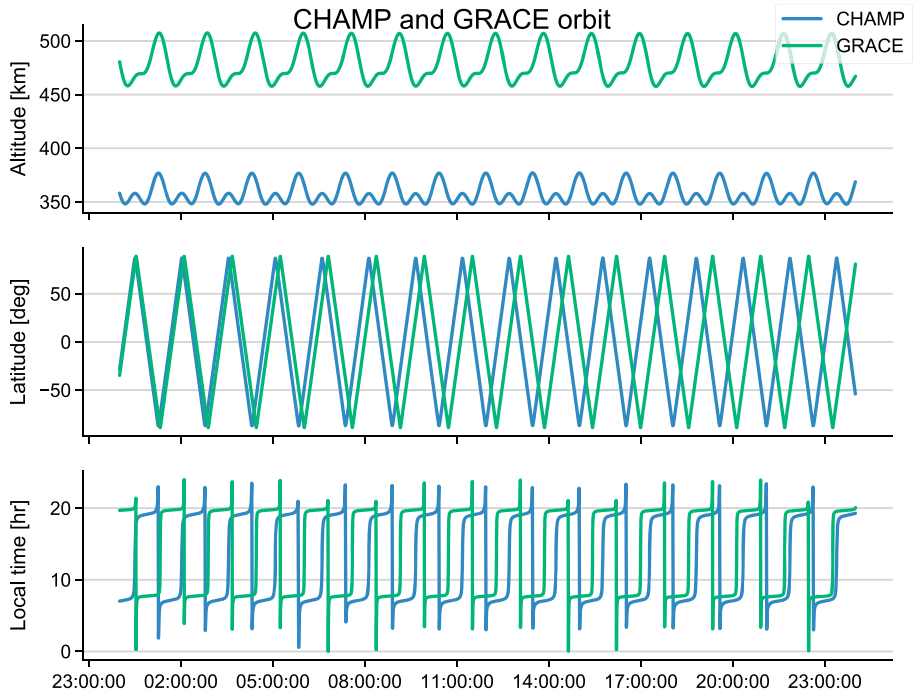

Figure 3. Shown are altitude, geographic latitude, and satellite local time for both the CHAMP and GRACE satellites. GRACE remains about 100 km above CHAMP at all times and both vary in location relative to each other throughout the day. The variable separation between the satellites is important to explore how assimilation of data from one satellite (CHAMP) influences neutral density specification at the location of the other (GRACE).

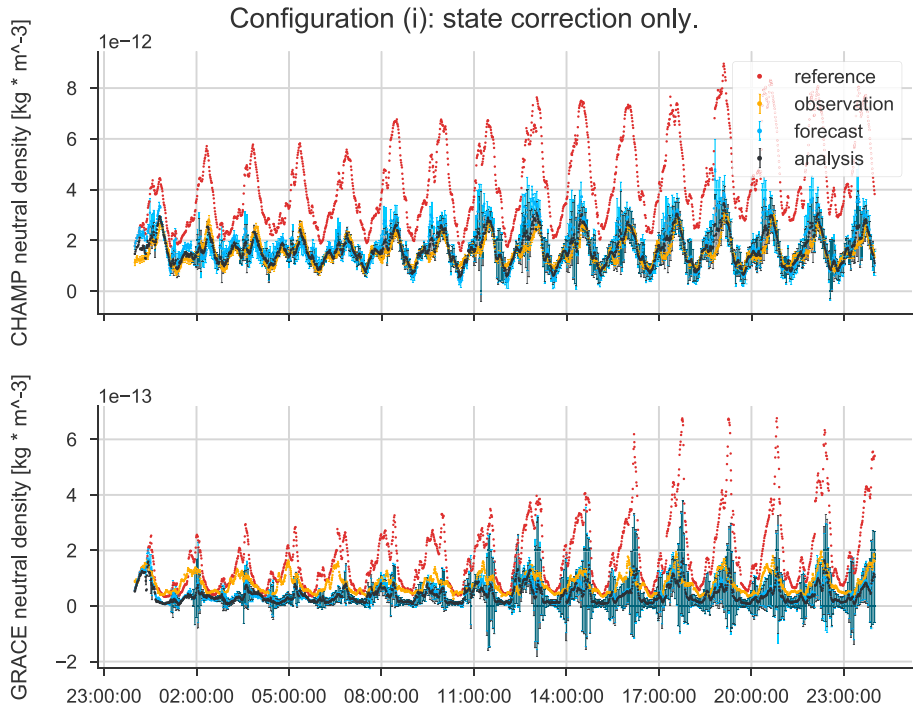

Figure 4. Neutral density results in Configuration (i): state correction. Top: assimilated observations and specification at the CHAMP satellite. Bottom: for comparison, observations, and specification at the GRACE satellite. Observations are shown in gold, forecast neutral density in blue, and final analysis specification in black. For reference, shown in red is the reference model run without assimilation. Error bars on gold observations show the uncertainty indicated by the data provider. Error bars on forecast and analysis points indicate the corresponding error estimates. Model specification comes into good agreement with CHAMP observations after three assimilation time steps and tracks observations for the remainder of the day. Bottom comparison illustrates a known bias between GRACE and CHAMP observations.

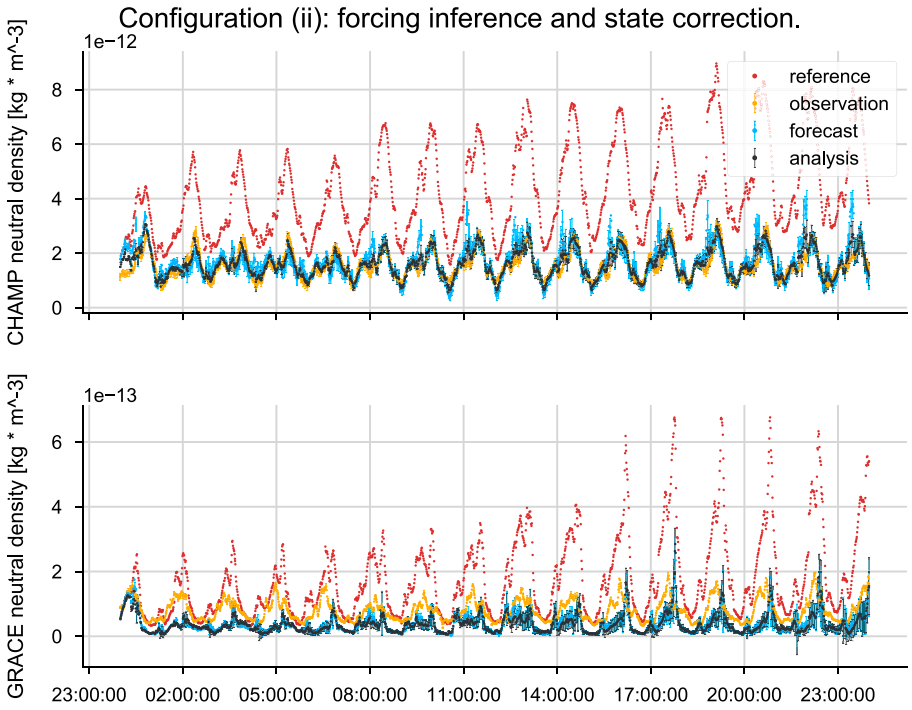

Figure 5. Neutral density results in Configuration (ii): both state correction and forcing estimation. Top: assimilated observations and specification at the location of CHAMP. Bottom: comparison with GRACE observations. Same format as Figure 4. In specification, forcing estimation in conjunction with state correction in Configuration (ii) performs slightly better than just state correction in Configuration (i); however, error bars are smaller indicating higher confidence in solution. Again, bottom plot shows known bias between GRACE and CHAMP observations.

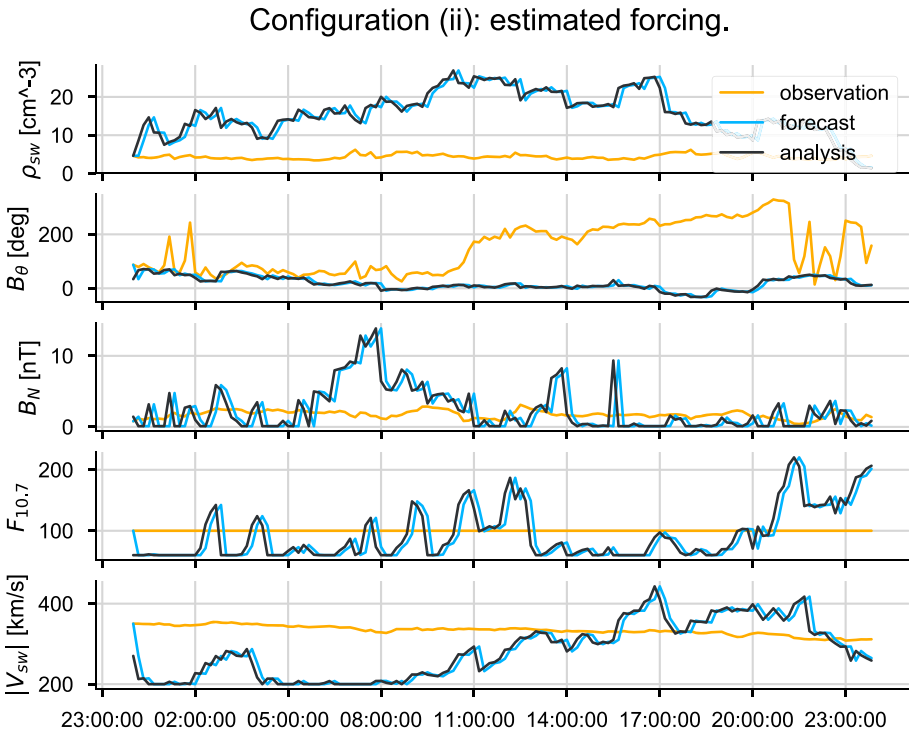

Figure 6. Estimated forcing resulting from Configuration (ii): both forcing inference and state correction. Observed forcing parameters shown in gold, forecast in blue, and final analysis estimate in black. The analysis estimate becomes the forecast value for the next assimilation time step. Note, the red reference in Figures 4, 5, and 7 are produced using the observations in gold. These estimated forcing parameters are not necessarily an estimate of reality; they are statistically determined to improve CTIPe agreement with observations.

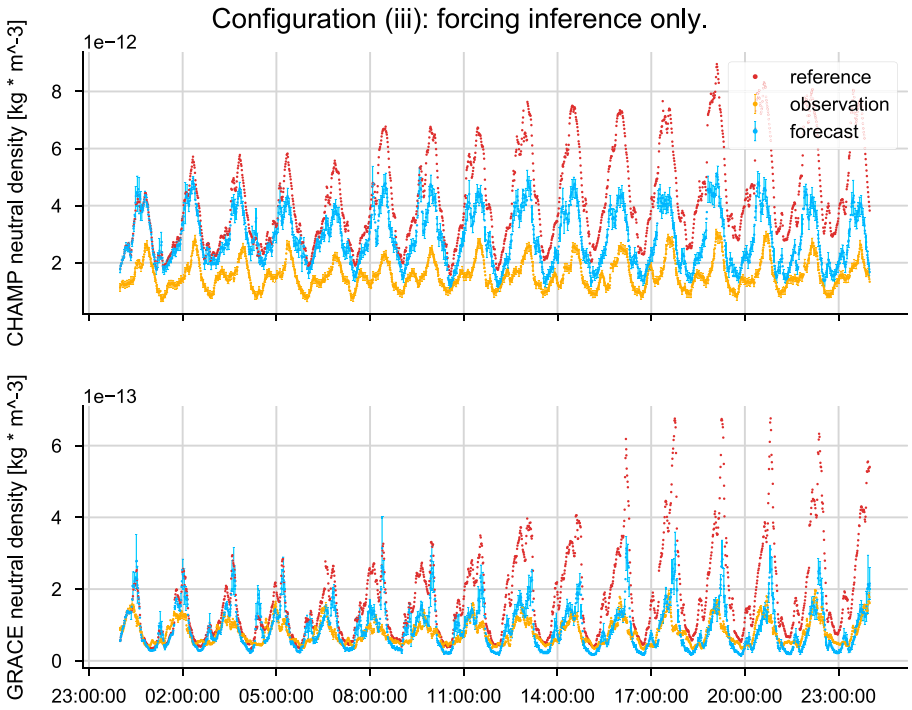

Figure 7. Neutral density results from Configuration (iii) only forcing estimation. Top: neutral density assimilated observations and specification at the CHAMP satellite. Bottom: for comparison, observations and specification at the GRACE satellite. Observations are in gold. Blue indicates the forecast. The analysis (black in previous plots) is not shown as this configuration does not modify the state and so Neutral Density specification is the same between forecast and analysis. The results indicate that the system can only react so quickly to certain changes in the forcing. It is conceivable that this configuration could effectively track the observations given a more accurate initial condition. As shown, the entire day is spent resolving the initial bias. The estimated forcing parameters improve agreement with observations relative to the reference.

forcing component is not ignored. For each assimilation time step, the estimated model forcing parameters of the special member are used for the special member over the next assimilation time step, and a perturbation of that for each ensemble member.

In Configuration (iii), only the model forcing parameters are estimated. The estimated forcing is handled in the same way as for Configuration (ii), the estimated forcing parameters for the special member are applied to the special member over the following assimilation time step, and a perturbation of that is used for each ensemble member. The model state component is completely ignored; it is not forecast by the background model. Instead, at each time step, each member retains their previous forecast state. While Configuration (iii) does not perform as well as Configuration (i) or (ii), it is included as a discussion point to highlight characteristics of the DA scheme.

5.2. Observations

Neutral density observations used in this experiment are derived from accelerometers flown aboard the polar orbiting Challenging Minisatellite Payload (CHAMP) (Reigber et al., 2002; Sutton et al., 2007) and Gravity Recovery and Climate Experiment (GRACE) (Sutton et al., 2007; Tapley et al., 2004) satellites. These accelerometers are sensitive enough to estimate in situ neutral density based on acceleration due to drag forces.

Figure 3 shows the orbits of the GRACE and CHAMP satellites. In all three configurations, the same observations from CHAMP are assimilated. Nominally, about 10 CHAMP measurements are available for each 10 min assimilation time step. Observations made by GRACE are not assimilated but are instead used for independent comparison at a point some distance away from CHAMP.

Observations of CTIPe model forcing parameters are provided by Space Weather Prediction Center. The $F_{10.7}$ index is measured by the Penticton Radio Observatory. The operational $F_{10.7}$ value used is the previous day averaged with the $F_{10.7}$ averaged over the previous 41 days (Codrescu et al., 2012). Solar wind parameters -

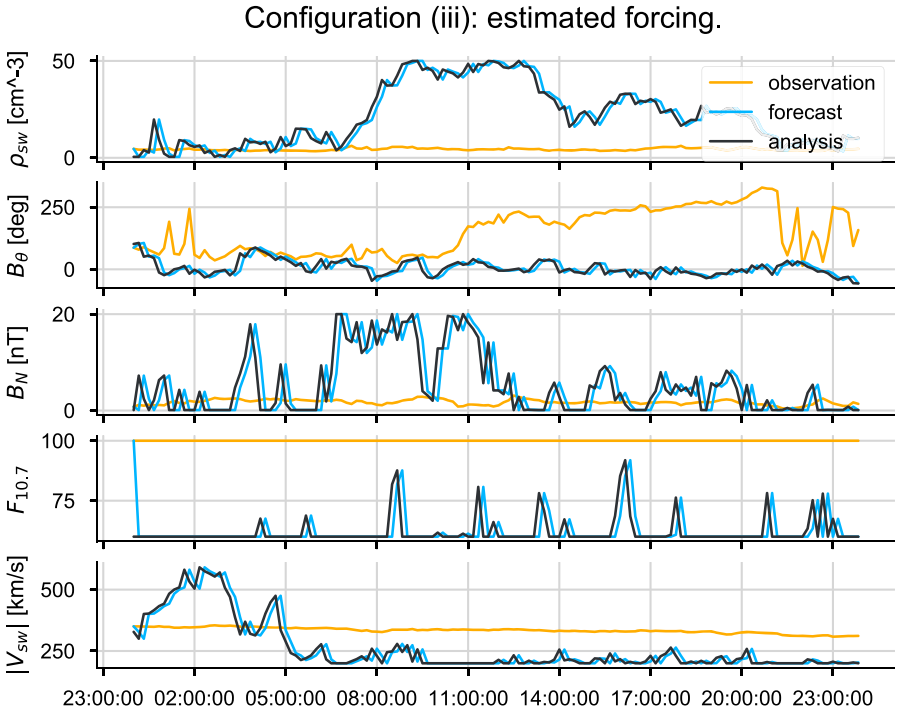

Figure 8. Estimated forcing parameters from Configuration (iii) only forcing estimation. Observed and estimated model input parameters shown in the same format as Figure 8. These estimated forcing parameters are definitely not an estimate of reality; instead, they reflect values chosen to improve model agreement with observations. Given the persistent bias between model state and observations and given the statistical ensemble correlations, physically impossible values may appear necessary to "force" the state, negative $F_{10.7}$, for example, but such results are clipped to a valid range. In this way, Configuration (iii) is limited in how quickly, for example, energy may be removed from the system.

$|V_{sw}|$, ρ_{sw}, B_N, and B_θ are observed by the Advanced Composition Explorer (ACE) satellite. Configuration (i) is forced according to these measurements, while Configuration (ii) and (iii) use these input measurements only for the initial assimilation time step and rely on estimated forcing parameters for subsequent assimilation time steps.

5.3. Results

Figure 4 shows the result of assimilating CHAMP neutral density measurements and comparison with GRACE observations via state correction only in Configuration (i). In all neutral density plots (Figures 4, 5, and 7), observations are shown in gold, blue indicates forecast neutral density specification before assimilation, and black indicates analysis specification after assimilation. Corresponding specification by the reference member is shown in red. Error bars on gold observations show the uncertainty indicated by the data provider. Error bars on forecast and analysis points indicate the corresponding forecast error and analysis error estimates.

Configuration (ii) produces the results shown in Figure 5. Estimation of forcing parameters combined with state correction produces the best agreement and lowest uncertainty results. Furthermore, this configuration most clearly illustrates the bias between GRACE and CHAMP observations. As Configuration (ii) also estimates the forcing parameters, these are shown in Figure 6. In both forcing parameter plots (Figures 6 and 8), observations are shown in gold, forecast in blue, and analysis in black. Note that the forecast values are simply the analysis values from the previous assimilation time step.

Finally, for Configuration (iii), the result of estimating only forcing parameters is shown in Figure 7. Since the state is not corrected via the Kalman update in this configuration, the analysis neutral density specification is the same as the forecast and is not shown. The corresponding Configuration (iii) estimated forcing parameters are displayed in Figure 8. While agreement with observations is worst in this configuration, it highlights both the response time to forcing parameters and points to reasons that estimating both state correction and forcing parameters in Configuration (ii) yields the best performance.

Table 1
Normalized Root Mean Square (NRMS) Difference Between Observations and Model for Each Configuration

	NRMS (x^f, y_{CHAMP})	NRMS (x^a, y_{CHAMP})	NRMS (x^f, y_{GRACE})	NRMS (x^a, y_{GRACE})
Configuration (i)	35%	16%	53%	56%
Configuration (ii)	30%	14%	58%	60%
Configuration (iii)	87%	-	57%	-

Note. Using the notation introduced in section 3, we present NRMS $(x, y) = \frac{\text{RMS}\,(y-h(x))}{\text{RMS}\,(y)}$, where RMS $(y) = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i)^2}$. The NRMS is presented for the forecast x^f, analysis x^a for both CHAMP and GRACE observations. In all Configurations, the reference is the same, with NRMS $(x_{\text{ref}}, y_{\text{CHAMP}}) = 179\%$ and NRMS $(x_{\text{ref}}, y_{\text{GRACE}}) = 170\%$.

The results of all three Configurations are summarized in Table 1 through presentation of a Normalized root-mean-square (NRMS) difference between the model specification and observations. In all Configurations, the results show significant improvement over the reference run.

6. Discussion and Conclusion

The results of each Configuration show a bias evident when comparing the agreement between assimilated CHAMP observations and CTIPe specification versus comparison with GRACE observations. To verify that the poor agreement in average between assimilated results and the comparison GRACE neutral density observations is the symptom of a systematic bias between the data sets, not an inability of the method to extrapolate to regions away from the ingested data set, the relative bias for neutral density specification between GRACE versus CTIPe and CHAMP versus CTIPe is calculated. The satellite observations were normalized to 400 km using the Mass-Spectrometer-Incoherent-Scatter (MSIS) model. For the full year of 2007, the relative bias has been estimated to be −0.08 between CHAMP and CTIPe, while for GRACE observations, the relative bias is −0.4. The conclusion is that there exists approximately 32% bias between GRACE and CHAMP observations. This agrees with a 30% bias calculated in Matsuo et al. (2012). For the remainder of the analysis, this bias will be ignored.

Configuration (i) is likely the most familiar in EnKF literature. It was presented first as a demonstration that the premise of an EnKF based on an ensemble distributed via randomized forcing can produce covariance statistic resulting in improvement of CTIPe global neutral density specification during quiet times through the assimilation of highly sparse measurements. This conclusion is reached by comparing the analysis result to the reference run, which indicates assimilation performs significantly better at reproducing observations than can be done using only the forcing parameters available operationally.

The improvement gained in Configuration (ii) in terms of slightly better specification and reduced uncertainty results from estimating forcing parameters which are then applied over the next assimilation time step. An effect of applying estimated forcing parameters is improved overall state consistency. The state of the model consists of more fields than just those included in the EnKF state vector. By estimating both the forcing and a correction to the state, elements not included in the filter are kept in better consistency via forcing with those that are corrected in place by the scheme.

It is important to highlight that the estimated forcing parameters do not represent an estimate of reality. The estimated parameters, for example, do not dispute solar wind observations; instead, they are statistically determined to improve agreement between CTIPe results and the assimilated observations. There are many reasons for the discrepancy between observed forcing parameters and estimated parameters. Examples include systematic biases and incorrectly tuned parameters within the model, which the scheme may attempt to compensate for. Furthermore, while the operational value of $F_{10.7}$ is measured once per day, the UV and EUV fluxes for which it is a proxy vary on much shorter time scales and can be estimated at a higher cadence. Finally, for solar wind values, which ultimately parameterize Joule heating through empirical models of auroral precipitation and convection electric fields, the Joule heating suggested by neutral density observations may be different than what is parameterized by the observed solar wind values. From the point of view of the model inputs, this is manifested in solar wind uncertainty but really represents uncertainty between

observation at ACE until Thermosphere-Ionosphere effectiveness, which includes propagation to Earth and furthermore through interaction with the magnetosphere. In summary, forcing parameters other than those observed may produce the best model agreement with observations.

A potential disadvantage of performing state correction as in Configurations (i) and (ii) is the disruption of state self-consistency. For example, the mass mixing ratio of the simulated atomic and molecular species must sum to 1. Unfortunately, the filter is currently unaware of this constraint and other physical limits within the system. To address this, the mixing ratios are normalized after assimilation to ensure they sum to 1; however, more investigation is necessary into the effects of assimilation on state self-consistency and how to address inconsistencies, should they occur.

Another disadvantage of state correction is the introduction of physically unsupported trajectories in state space—consider that correcting the model likely produces a state space trajectory that cannot be realized through physical dynamics of the system, or one that is not implemented in the model. In fact, this is the fundamental way in which results from Configurations (i) and (ii) are different from Configuration (iii). In practice, this appears to be of little consequence to the goal of improving specification but could become detrimental to prediction. Careful examination of the state space trajectory under assimilation versus without could even shed light on missing physics.

Configuration (iii) represents one possible solution to prevent nonphysicalities in state space. Estimation of only the forcing parameters restricts the scheme to steer the forecast model according only to the dynamics implemented. As the results indicate, the improvement is not as significant compared to the other configurations. At first, the result that state correction fares better than pure forcing correction may appear to contradict assertions that the system is strongly forced. This, however, is not a contradiction but is completely a result of the system response time to changes in the forcing parameters. Consider, for example, that it is easy to rapidly inject energy into the system through a combination of forcing parameters but hard to rapidly remove energy. State correction in Configurations (i) and (ii) fares better because we are able to impose a state with less energy, instead of relying only on natural processes in the case of Configuration (iii) to remove the energy.

Physical limits to the response time of the system are apparent in another way in Configuration (iii) in that the estimated forcing is frequently clipped to a range between empirical nominal minimum and maximum observed values. These bounds are necessitated by the filter returning physically impossible values at times. According to the statistics of the ensemble, physically impossible values that would crash the model may be necessary to move the state to the measurements. Of course, these values cannot be allowed. We suspect that if the state were initially very close to reproducing observations, more nominal forcing values would be estimated and the boundaries on these parameters would not be encountered. In such case, we would hope to investigate how such estimated parameters compared to observed values.

Similar to other work (Chen et al., 2016; Scherliess et al., 2004), we find good performance with assimilation time steps on the order of 10 to 30 min. Preliminary experiments (not shown) indicate deteriorating performance with assimilation time steps under 5 min. Between 10 and 30 min, there appears to be a trade-off between forecast performance and analysis performance—dependent on Configuration as well. For the sake of consistency between configurations presented, each uses 10 min assimilation time steps; however, interestingly Configuration (iii) appears to produce better results with 30 min assimilation time steps. A more detailed investigation of the configuration space is necessary in this case.

Unlike virtually all other Thermosphere-Ionosphere assimilation schemes in the literature (Chartier et al., 2016; Godinez et al., 2015; Hsu et al., 2014; Matsuo et al., 2013; Morozov et al., 2013; Solomentsev et al., 2012), we do not localize the covariance matrix. In these previous studies, localization is applied to minimize the effect of spurious correlations. Spurious correlations are an artifact of the initial randomization of a large number of state elements across a small ensemble. Consider that the size of the EnKF state vector is usually many orders of magnitude greater than the number of ensemble members. Given such proportions, it is impossible to generate an ensemble without giving rise to spurious correlations. In our assimilation scheme, the ensemble distribution is created through the randomization of six forcing elements across a 75 member ensemble. It becomes improbable for spurious correlations to arise when the size of the ensemble is an order of magnitude larger than the number of randomized elements. We speculate that since the ensemble uncertainty is distributed in state space along the few degrees of freedom of the forcing, the covariances that arise are global

and tied in correlation with the forcing parameters. In other words, the six forcing parameters form a basis for the uncertainty over which our 75 member ensemble is able to effectively sample. Resulting corrections due to observations are made globally through the entire state.

To conclude, we have explored initial results demonstrating that our Ensemble Kalman Filter-based Data Assimilation scheme is capable of improving global neutral density specification based on highly sparse measurements during geomagnetically quiet conditions at solar minimum. The key to this result is recognition that the primary system uncertainty lies in the system forcing—and that consequently ensemble members should be distributed by randomizing the forcing applied to each. It appears that, at least for these conditions, without randomizing the initial condition, the need for localization is eliminated and neutral density measurements can be utilized to apply global neutral density corrections. With the current implementation in place, we anticipate turning our attention to assimilation of measurements during storm conditions and extending the assimilation to include ionospheric measurements.

Acknowledgments
The authors thank Eric Sutton for providing the CHAMP and GRACE measurements used in this study. We also thank Karen O'Loughlin for constructive criticism and assistance in preparation for publication as well as Tim Fuller-Rowler for helpful conversations. We are grateful to the anonymous reviewers for feedback that has strengthened and clarified this work. Data results are available at https://doi.org/10.5281/zenodo.1146765.

References

Akmaev, R. A. (2011). Whole atmosphere modeling: Connecting terrestrial and space weather. *Reviews of Geophysics*, *49*, RG4004. https://doi.org/10.1029/2011RG000364

Chartier, A. T., Matsuo, T., Anderson, J. L., Collins, N., Hoar, T. J., Lu, G., … Bust, G. S. (2016). Ionospheric data assimilation and forecasting during storms. *Journal of Geophysical Research: Space Physics*, *121*, 764–778. https://doi.org/10.1002/2014JA020799

Chen, C. H., Lin, C. H., Matsuo, T., Chen, W. H., Lee, I. T., Liu, J. Y., … Hsu, C. T. (2016). Ionospheric data assimilation with thermosphere-ionosphere-electrodynamics general circulation model and GPS-TEC during geomagnetic storm conditions. *Journal of Geophysical Research: Space Physics*, *121*, 5708–5722. https://doi.org/10.1002/2015JA021787

Codrescu, M. V., Fuller-Rowell, T. J., & Kutiev, I. S. (1997). Modeling the *F* layer during specific geomagnetic storms. *Journal of Geophysical Research*, *102*, 14,315–14,329. https://doi.org/10.1029/97JA00638

Codrescu, M. V., Fuller-Rowell, T. J., & Minter, C. F. (2004). An ensemble-type Kalman filter for neutral thermospheric composition during geomagnetic storms. *Space Weather*, *2*, S11003. https://doi.org/10.1029/2004SW000088

Codrescu, M. V., Fuller-Rowell, T. J., Munteanu, V., Minter, C. F., & Millward, G. H. (2008). Validation of the coupled thermosphere ionosphere plasmasphere electrodynamics model: CTIPE-mass spectrometer incoherent scatter temperature comparison. *Space Weather*, *6*, S09005. https://doi.org/10.1029/2007SW000364

Codrescu, M. V., Negrea, C., Fedrizzi, M., Fuller-Rowell, T. J., Dobin, A., Jakowsky, N., … Maruyama, N. (2012). A real-time run of the Coupled Thermosphere Ionosphere Plasmasphere Electrodynamics (CTIPe) model. *Space Weather*, *10*, S02001. https://doi.org/10.1029/2011SW000736

Crowley, G., Emery, B. A., Roble, R. G., Carlson, H. C., Salah, J. E., Wickwar, V. B., … Marcos, F. A. (1989). Thermospheric dynamics during September 18–19, 1984: 2. Validation of the NCAR Thermospheric General Circulation Model. *Journal of Geophysical Research*, *94*(A12), 16,945–16,959. https://doi.org/10.1029/JA094iA12p16945

Crowley, G., Hackert, C. L., Meier, R. R., Strickland, D. J., Paxton, L. J., Pi, X., … Wene, G. (2006). Global thermosphere-ionosphere response to onset of 20 November 2003 magnetic storm. *Journal of Geophysical Research*, *111*, A10S18. https://doi.org/10.1029/2005JA011518

Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, *99*(C5), 10,143–10,162. https://doi.org/10.1029/94JC00572

Fedrizzi, M., Fuller-Rowell, T. J., & Codrescu, M. V. (2012). Global Joule heating index derived from thermospheric density physics-based modeling and observations. *Space Weather*, *10*, S03001. https://doi.org/10.1029/2011SW000724

Fuller-Rowell, T. J., Akmaev, R. A., Wu, F., Anghel, A., Maruyama, N., Anderson, D. N., … Millward, G. (2008). Impact of terrestrial weather on the upper atmosphere. *Geophysical Research Letters*, *35*, L09808. https://doi.org/10.1029/2007GL032911

Fuller-Rowell, T. J., Codrescu, M. C., & Wilkinson, P. (2000). Quantitative modeling of the ionospheric response to geomagnetic activity. *Annales Geophysicae*, *18*(7), 766–781. https://doi.org/10.1007/s00585-000-0766-7

Fuller-Rowell, T. J., & Evans, D. S. (1987). Height-integrated Pedersen and Hall conductivity patterns inferred from the TIROS-NOAA satellite data. *Journal of Geophysical Research*, *92*(A7), 7606–7618. https://doi.org/10.1029/JA092iA07p07606

Fuller-Rowell, T. J., & Rees, D. (1980). A three-dimensional time-dependent global model of the thermosphere. *Journal of the Atmospheric Sciences*, *37*(11), 2545–2567. https://doi.org/10.1175/1520-0469(1980)037<2545:ATDTDG>2.0.CO;2

Gillijns, S., Mendoza, O. B., Chandrasekar, J., De Moor, B., Bernstein, D. S., & Ridley, A. (2006). What is the ensemble Kalman filter and how well does it work, *IEEE 2006 American Control Conference* (pp. 6). Minneapolis, MN: IEEE. https://doi.org/10.1109/ACC.2006.1657419

Godinez, H. C., Lawrence, E., Higdon, D., Ridley, A., Koller, J., & Klimenko, A. (2015). Specification of the Ionosphere-Thermosphere Using the Ensemble Kalman Filter. In S. Ravela & A. Sandu (Eds.), *Dynamic Data-Driven Environmental Systems Science, Lecture Notes in Computer Science* (Vol. 8964, pp. 274–283). Cham: Springer. https://doi.org/10.1007/978-3-319-25138-7_25

Guennebaud, G., Jacob, B., & Lenz, M. (2010). Eigen v3. Retrieved from https://doi.org/http://eigen.tuxfamily.org

Hinteregger, H. E., Fukui, K., & Gilson, B. R. (1981). Observational, reference and model data on solar EUV, from measurements on AE-E. *Geophysical Research Letters*, *8*(11), 1147–1150. https://doi.org/10.1029/GL008i011p01147

Hsu, C.-T., Matsuo, T., Wang, W., & Liu, J.-Y. (2014). Effects of inferring unobserved thermospheric and ionospheric state variables by using an Ensemble Kalman Filter on global ionospheric specification and forecasting. *Journal of Geophysical Research: Space Physics*, *119*, 9256–9267. https://doi.org/10.1002/2014JA020390

Huba, J. D., Schunk, R. W., & Khazanov, G. V. (2014). *Modeling the ionosphere-thermosphere*. Chichester, UK: John Wiley.

Johnson, R. M., & Killeen, T. L. (1995), *The upper mesosphere and lower thermosphere: A review of experiment and theory, Geophysical Monograph Series* (pp. 356). Washington, DC: American Geophysical Union.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, *82*(1), 35–45. https://doi.org/10.1115/1.3662552

KORENKOV (2013). *Ionospheric modeling*. Birkhäuser Basel: Springer Basel AG.

Matsuo, T., Fedrizzi, M., Fuller-Rowell, T. J., & Codrescu, M. V. (2012). Data assimilation of thermospheric mass density. *Space Weather*, *10*, S05002. https://doi.org/10.1029/2012SW000773

Matsuo, T., Lee, I.-T., & Anderson, J. L. (2013). Thermospheric mass density specification using an ensemble Kalman filter. *Journal of Geophysical Research: Space Physics*, *118*, 1339–1350. https://doi.org/10.1002/jgra.50162

Millward, G. H., Moffett, R. J., Quegan, S., & Fuller-Rowell, T. J. (1996). A coupled thermosphere-ionosphere-plasmasphere model (CTIP). In R. W. Schunk (Ed.), *STEP handbook on ionospheric models* (pp. 239–279). Logan, UT. Utah State University.

Moe, K., & Moe, M. M. (2011). Operational models and drag-derived density trends in the thermosphere. *Space Weather*, *9*, S00E10. https://doi.org/10.1029/2010SW000650

Morozov, A. V., Ridley, A. J., Bernstein, D. S., Collins, N., Hoar, T. J., & Anderson, J. L. (2013). Data assimilation and driver estimation for the global ionosphere-thermosphere model using the ensemble adjustment Kalman filter. *Journal of Atmospheric and Solar-Terrestrial Physics*, *104*, 126–136. https://doi.org/10.1016/j.jastp.2013.08.016

Reigber, C., Lühr, H., & Schwintzer, P. (2002). CHAMP mission status. *Advances in Space Research*, *30*(2), 129–134. https://doi.org/10.1016/S0273-1177(02)00276-4

Scherliess, L., Schunk, R. W., Sojka, J. J., & Thompson, D. C. (2004). Development of a physics-based reduced state Kalman filter for the ionosphere. *Radio Science*, *39*, RS1S04. https://doi.org/10.1029/2002RS002797

Solomentsev, D. V., Khattatov, B. V., Codrescu, M. V., Titov, A. A., Yudin, V., & Khattatov, V. U. (2012). Ionosphere state and parameter estimation using the ensemble square root filter and the global three-dimensional first-principle model. *Space Weather*, *10*, S07004. https://doi.org/10.1029/2012SW000777

Sutton, E. K., Nerem, R. S., & Forbes, J. M. (2007). Density and winds in the thermosphere deduced from accelerometer data. *Journal of Spacecraft and Rockets*, *44*(6), 1210–1219.

Tapley, B. D., Bettadpur, S., Watkins, M., & Reigber, C. (2004). The gravity recovery and climate experiment: Mission overview and early results. *Geophysical Research Letters*, *31*, L09607. https://doi.org/10.1029/2004GL019920

Weimer, D. R. (2005). Improved ionospheric electrodynamic models and application to calculating Joule heating rates. *Journal of Geophysical Research*, *110*, A05306. https://doi.org/10.1029/2004JA010884