

Qualifying and Reducing Neutral Density Uncertainty for Precise Orbit Determination using Physics-Based Data Assimilations

Nicholas Dietrich

University of Colorado Boulder

Tomoko Matsuo¹, Chih-Ting Hsu²

¹University of Colorado Boulder, ²High Altitude Observatory, NCAR

ABSTRACT

With increasing space traffic in low Earth orbit (LEO), it is critical to better quantify uncertainties arising from neutral densities and satellite drag that impact orbit propagation. These uncertainties are largely due to incomplete knowledge on the atmosphere's response to geomagnetic storms and solar flares, as well as on atmospheric temperature and compositions. The lack of composition information is especially true for helium that begins to dominate in LEO altitudes, as they affect the ballistic coefficient estimation and mean mass density. In-situ observations of neutral densities are limited only along orbital tracks, and empirical density models often used in orbit determination struggle to predict densities during space weather events. Additionally, variability of helium compositions is poorly understood in the upper atmospheric science community with virtually no direct observations of atmospheric compositions. In this study we use a first-principles, physics-based model, Thermosphere Ionosphere General Circulation Model (TIEGCM), to globally estimate and predict upper atmospheric states. In TIEGCM, solar irradiance is parameterized by the F10.7 index, and the intensity of geomagnetic storm drivers is represented by the Kp index. Using an ensemble square root filter that incorporates the strong coupling between the thermosphere and ionosphere into both the measurement and time update steps of the filter, indirect but abundant radio occultation electron density observations can be assimilated into TIEGCM to estimate mass densities. Previous work have shown that neutral density estimation can be improved by this approach, however, the helium composition estimation and uncertainty quantification have been overlooked. This work shows for the first time that neutral densities and helium compositions can be better estimated as a result of electron density data assimilation. An additional focus of this work is the use of the particle filter to estimate the F10.7 and Kp drivers to reduce TIEGCM bias and estimate probabilistic driver distributions for uncertainty quantification. Here TIEGCM is used as a time-dependent forward model, mapping driver information to neutral densities so that the density observations can be used for driver estimation. Results show a reduction in model bias, and support the potential for the framework to estimate driver uncertainties. With promising results, both ensemble and particle methods are implemented with the intention of performing end-to-end uncertainty quantification, propagating uncertainties in atmospheric compositions and drivers to orbit position errors.

1. INTRODUCTION

Atmospheric neutral density variability poses significant issues in estimating satellite forecasting uncertainties in low Earth orbit (LEO). Satellite drag remains the largest source of uncertainty for LEO satellites, and is serious hurdle in precisely managing the all the satellites in the LEO population. Reducing satellite position uncertainties will be necessary in keeping the number of conjunction alerts within a manageable level [3]. During geomagnetic storms and solar flares, neutral density becomes highly variable, causing satellite position uncertainty ellipsoids to expand and can result in untenable numbers of conjunction alerts. Specifying neutral density ρ in the satellite drag equation is the focus of this paper

$$\mathbf{a} = \frac{1}{2} C_D \frac{A}{M} \rho v \mathbf{v} \quad (1)$$

It is difficult to gain global understanding of neutral densities as observations are only available through on-board accelerometers and precise orbit determination, limiting data to along orbit tracks. Additionally, first-principles, physics-based models of the upper atmosphere have advantages over current neutral density empirical models, like HASDM

[17], in that they can better model and forecast global states during geomagnetic storm events. For these two reasons, we turn to physics-based models and in-direct observations to estimate and forecast the thermosphere during space weather events.

This paper discusses two methods for specifying neutral densities. The first method is to assimilate radio occultation (RO) observations for electron densities into a physics-based model, and use the strong coupling between the thermosphere and ionosphere to provide neutral state updates. Previous work has shown RO can be effective in improving neutral density estimation [11], and this work will look to build upon those results by assessing RO observations' ability to specify helium compositions. The second method is the development of a particle filter to estimate model drivers. The particle filter will use real observations to reduce model biases and estimate the driver distribution to produce a time-varying probability distribution of solar and magnetospheric drivers. This is implemented with the goal of achieving more accurately represented atmospheric uncertainties. Both these methods will enable the eventual goal of performing end-to-end uncertainty quantification. Beginning at one end with estimated driver distributions, and then mapping these uncertainties through a physics-based model, we can better quantify neutral density uncertainties. Improved neutral density uncertainties may then be used to better represent orbit position error ellipsoids. This end-to-end uncertainty quantification will enable the accurate quantification of orbit errors during storm times.

1.1 Thermosphere-Ionosphere Coupling

To provide the physics-based coupling needed to specify the relationship between the thermosphere and ionosphere, the Thermosphere Ionosphere Electrodynamics General Circulation Model (TIEGCM), developed by NCAR [13], [15], is used. TIEGCM is a first-principles model that extends from 90 km to 600 km, encompassing much of the relevant LEO satellite altitudes. TIEGCM is driven by the two major upper atmospheric drivers: solar forcing through extreme ultraviolet (EUV) radiation and magnetospheric forcing through joule heating and particle precipitation. EUV radiation produces a rapid response to Earth's thermosphere on Earth's day-side, ionizing and heating neutral gases [10]. Joule heating and particle precipitation occurs at the high latitude and arrive along Earth's magnetic field lines, and has a more delayed response in impacting lower latitudes. These two driver parameters are controlled in TIEGCM using the F10.7 and Kp indices. The F10.7 index is a daily measurement of the 10.7 cm wavelength of solar radio flux and functions as a proxy for EUV radiation. The Kp index is a logarithmic scale for measuring geomagnetic activity, with values 4 and greater considered a geomagnetic storm event. Magnetospheric forcings may also be specified through hemispheric power (HP) and cross-tail potential drop (Φ). As these drivers function as proxies for more complicated processes, using their true observed quantities may result in considerable model biases. Unlike the lower atmosphere which is a highly chaotic system dependent on initial conditions, the upper atmosphere (above 70 km) is quite strongly controlled by the external drivers derived from the Sun.

Species compositions for atomic oxygen (O), molecular oxygen (O₂), nitrogen (N₂), and helium (He) and consequently neutral density may be described using scale height H seen in Eqs. 2 and 3.

$$H = \frac{k_B T}{\bar{m}g} = \frac{RT}{\bar{M}g} \quad (2)$$

$$\rho = \rho_0 \exp\left(-\frac{alt}{H}\right) \quad (3)$$

where R [J/K/kmol] is the universal gas constant, g [m/s²] is the gravitational constant, \bar{M} [kg/kmol] is the mean molecular mass and T [K] is temperature.

1.2 Observation Data

RO remote sensing has proved to be an incredible tool for understanding Earth's ionosphere. Using the radio signals from already existing Global Navigation Satellite Systems (GNSS) such as USA's GPS, Russia's GLONASS, China's BeiDou and the EU's Galileo, satellites in LEO can measure the ionosphere's effects on these signals and produce vertical electron density profiles (EDPs) of the ionosphere. With the existing global coverage of the GPS constellation, the COSMIC constellations produce RO observations that span the globe. The currently operating COSMIC II constellation can yield RO soundings with GLONASS and Galileo systems in addition to GPS. The global coverage of RO data and fast data latency make these appealing data sources for space weather prediction.

For information on true thermospheric states, derived neutral densities from CHAMP are utilized for driver estimation. These densities are derived by Eric Sutton using CHAMP's on-board accelerometer [19]. CHAMP has a polar orbits around altitudes of 340 km, providing good coverage of the day and night sides of the atmosphere at LEO altitudes.

2. HELIUM ESTIMATION

Global helium compositions are a dark spot for upper atmospheric science. While global composition information is available for the column integrated O/N₂ composition ratio, helium is not well observed. Helium compositions observations have only come from limited mass spectrometer measurements or been derived from precise orbit determination [21], with most direct helium observations coming from satellites flown in the 60s and 70s [7], [14], [5]. Helium compositions become dominate at higher LEO altitudes, but can reach lower altitudes through the phenomenon known as the winter bulge first discovered in [7]. When temperatures become small, scale heights decrease and the atmosphere contracts, causing lighter species like helium become the dominant species at lower altitudes.

Helium uncertainty introduces errors into three major satellite drag components. The first being the mean mass density calculation, with helium possessing a fourth of the mass of atomic oxygen. The second is in the calculation of scale height, which additionally depends on calculating the mean molecular mass. These scale height equations are needed to find densities at the necessary altitudes. Finally, composition knowledge impacts the drag coefficient C_D that rely on gas-surface interaction models to calculation. Without Helium knowledge, unforeseen uncertainties are introduced into orbit determination.

To help illustrate the impact helium uncertainty can cause, a polar orbit was simulated and propagated for 72 hours. For neutral densities at orbit locations, two TIEGCM runs are used with quiet solar conditions: one with Helium and one without. In-track orbit errors are calculated after propagating the satellite for a given amount of time, as seen in Table 1. While orbit errors are generally larger than those at lower altitudes, errors can increase to up to 1 km after 72 hours of propagation.

Table 1: Polar In-track Orbit Errors Using TIEGCM Fields With and Without Helium

Altitude	C_D	12 Hour Error	24 Hour Error	48 Hour Error	72 Hour Error
400 km	2.2	0.013 km	0.057 km	0.135 km	0.552 km
400 km	4	0.024 km	0.105 km	0.244 km	1.005 km
500 km	2.2	0.013 km	0.054 km	0.120 km	0.492 km
500 km	4	0.023 km	0.097 km	0.218 km	0.895 km

To address this issue, alternative means for estimating helium are needed to improve knowledge of the LEO environment. Previous work from [11] has shown COSMIC RO data assimilated into TIEGCM has shown to improve state updates for neutral temperature, and consequently neutral density. With an updated version of TIEGCM with helium implemented [20], these experiments can be replicated to assess their ability to improve helium estimates.

2.1 Ensemble Adjustment Kalman Filter

Ensemble Kalman Filter (EnKF) is an extension of base Kalman filter Bayesian methods for high dimensional systems [1]. With atmospheric model states commonly around the order of $O(n) = 10^3 - 10^8$, the traditional linear Kalman filter can not handle matrix inverse of this size. To operate on these high-dimensional systems, the EnKF relies on sample statistics, using a low-rank approximation of the covariance matrix. The covariance subspace is approximated by subspace of size q , with $q \ll n$, to generate covariance matrices of rank q . The covariance \mathbf{P} is approximated by

$$\mathbf{P} \approx DD^T \quad (4)$$

where D are the data matrices constructed using ensemble members

$$D = \frac{1}{\sqrt{q-1}} \begin{bmatrix} (\mathbf{x}^{(1)} - \boldsymbol{\mu}) & (\mathbf{x}^{(2)} - \boldsymbol{\mu}) & \dots & (\mathbf{x}^{(q)} - \boldsymbol{\mu}) \end{bmatrix} \in \mathbb{R}^{n \times q} \quad (5)$$

where $\boldsymbol{\mu} \in \mathbb{R}^n$ is the ensemble state mean. Commonly, around $q = 80 - 100$ ensembles are used to approximate the sample statistics. In implementation, the covariance matrix does not need to be calculated.

Like the base Kalman filter, the EnKF uses a recursive algorithm with a forecast step and a measurement update step. For the forecast step, the full non-linear model dynamics, F , are taken advantage and all ensemble members are propagated forward from the current time t_i to the next observation time t_{i+1} .

$$\mathbf{x}^{fk}(t_{i+1}) = F\left(\mathbf{x}^{ak}(t_i), \mathbf{d}(t_i)\right) \quad k = 1, \dots, q \quad (6)$$

where k is the index of the ensemble member, $\mathbf{x}^k(t_i)$ is current analysis or posterior state estimate, $\mathbf{x}^k(t_{i+1})$ is the forecast or prior state at the next time step and $\mathbf{d}(t_i)$ are model drivers.

For the measurement update, the Ensemble Adjustment Kalman Filter (EAKF) detailed in [1] and [2] and implemented by the NCAR Data Assimilation Research Testbed (DART). For the rest of this paper, the ensemble filter will be referenced as the EAKF. This is a deterministic ensemble update method that assumes independent observations, allowing measurement updates to be performed sequentially. Model states are first mapped from state space to observation space using the non-linear forward model h

$$\mathbf{y}^{fk}(t_{i+1}) = h(\mathbf{x}^{fk}(t_{i+1})) \quad (7)$$

For individual observations j , states are then sequentially updated in observation space, and the observation space increments are mapped back to state space

$$\left(\mathbf{x}^k(t_{i+1})\right)_j = \mathbf{x}^{fk}(t_{i+1}) + \alpha \frac{(D_x)_j D_{y_j}^T}{D_{y_j} D_{y_j}^T} \left(\mathbf{y}_j^k(t_{i+1}) - \mathbf{y}_j^{fk}(t_{i+1})\right) \quad (8)$$

for each ensemble $k = 1, \dots, q$, using state data matrix D_x and observation data matrix D_y . α is a localization parameter multiplied by the Shur product. This update step is repeated for each observation, $j = 1, \dots, m$ where the state data matrix $(D_x)_j$ is updated each observation using the newly calculated $(\mathbf{x}^k(t_{i+1}))_j$.

For the EAKF, it is not necessary that observation states be directly related to updated model states through h . In the traditional Kalman filter, a linearized forward model is needed to map updates in observation updates to updates in state space. However, in Eq. 8, the mapping is performed through a linear regression coefficient. This linear regression coefficient is constructed using ensemble-based covariance information, relying on correlations between model states and observations. For this experiment, observations are COSMIC electron densities, and updated states are neutral states like temperature, molecular compositions, and neutral winds.

2.2 Experiments

To assess the ability of RO electron densities to improve helium estimates, an observing system simulation experiment (OSSE) is used. In an OSSE, a model is run with given driver conditions and is treated as the "true" atmosphere. From this run, real COSMIC RO measurement locations are sampled from in the model and perturbed using the corresponding observation variance. With a synthetic set of observation data, these data are assimilated into an ensemble run using the EAKF implemented by DART. With an OSSE, estimated atmospheric states may be directly verified against the true atmosphere to assess filter performance. For these helium estimation experiments, the true model was run using the drivers $\theta_{F107} = 74$, $\theta_{HP} = 55$ and $\theta_{\Phi} = 18$ and sampled for synthetic observations using $\varepsilon \sim N(0, \mathbf{R})$. A 90 member ensemble was initialized through sampling from a Gaussian distribution of drivers with means biased from the true data. These sample driver distributions had mean and variances $\theta_{F107} \sim N(69, 5^2)$, $\theta_{HP} \sim N(45, 10^2)$ and $\theta_{\Phi} \sim N(16, 2^2)$. All ensembles are initialized with these sampled drivers and spun-up for 15 days, allowing the atmosphere to reach steady state. The filter is run for two days, from UT 0 on June 23rd, 2008 to UT 22 on June 24th, 2008. The prior ensembles begin at June 23rd UTC00, and RO data is assimilated at each subsequent hour. Four experiments are performed using the same set of initialized ensembles and altering which neutral states were updated, shown in Table 2. Updated states include electron density \mathbf{f}_{e^-} , O^+ ion density \mathbf{f}_{O^+} , neutral temperature \mathbf{f}_T , helium mixing ratio \mathbf{f}_{He} , and neutral winds \mathbf{f}_{VN} and \mathbf{f}_{UN} in the north-south, .

Table 2: OSSE States Updated using COSMIC RO EDPs

Experiment	Total Number of Updated States	\mathbf{x}
1	225,504	$[\mathbf{f}_{e^-}; \mathbf{f}_{O^+}; \mathbf{f}_T]$
2	300,672	$[\mathbf{f}_{e^-}; \mathbf{f}_{O^+}; \mathbf{f}_T; \mathbf{f}_{He}]$
3	375,840	$[\mathbf{f}_{e^-}; \mathbf{f}_{O^+}; \mathbf{f}_T; \mathbf{f}_{VN}; \mathbf{f}_{UN}]$
4	300,672	$[\mathbf{f}_{e^-}; \mathbf{f}_{O^+}; \mathbf{f}_{VN}; \mathbf{f}_{UN}]$

2.3 Experiment Results

To show these results, global root mean square error (RMSE) is calculated each hour using the error between the prior experiment states and the truth run. RMSEs are calculated through areas of equal pressure levels, the unit TIEGCM

uses as its vertical axis. The top and bottom pressure levels of TIEGCM are excluded. One pressure level can contain a wide span of altitudes, so the top few pressure levels are focused on as these best span LEO altitudes and are where helium becomes abundant, about 360 to 530 km. Additionally, due to density varying exponentially with altitude, focusing on only the top few pressure levels will accurately represent the error at high altitudes. RMSEs are plotted for all four experiments as well as the control which had no RO data assimilation.

Examining Fig. 1 that shows neutral density RMSE, there is an immediate improvement over the control in all the experiments that update temperature. Errors reach a minimum at hour 13, and then increase and decrease in a diurnal cycle, peaking 12 hours later and reaching a minimum 24 hours later. This improvement is attributed to improved temperature estimation which is the largest controlling factor in calculating neutral density. These results are consistent with what was seen in [11] that showed improvements in temperature when directly updated with COSMIC RO. As for updating helium directly in experiment 2, there is slightly improvement over only experiment 1, however beginning around UT 12, the error begins to diverge and drastically increase. Updating neutral winds directly in experiment 3 looks to offer some improvement, however there is negligible impact for these time scales. Updating purely neural winds, seen in experiment 4, provides no impact on neutral densities.




Fig. 1: Global neutral density RMSE spanning the top three TIEGCM pressure levels, altitudes 360 km ~ 530 km.

The new question this study looks to answer is whether COSMIC RO can provide improvement to helium estimation. Analyzing Fig. 2, there is improved errors for experiments 2 and 3 over the control. These errors improve at a slower rate than density, suggesting a lag time in helium estimation improvement. Helium errors reach a minimum at UT 4 on June 24th. Like in Fig. 1, updating helium directly looks to initially provide benefit for the first few time DA cycles. However, around hour 7 the the error grows and increases beyond the control. This suggests some potential for direct helium updates, however, these results will need to be further analyzed to evaluate its usefulness. Experiment 3 shows updating neutral winds can have improve helium estimation slightly over only updating temperature, moreso than was shown for density RMSE. This is confirmed in experiment 4, that while there was no impact on neutral densities, updating neutral winds can improve helium estimation. A diurnal cycle is additionally visible in experiments 1, 3 and 4.

To better understand why helium composition is improved via temperature estimation, we can look at scale height RMSE shown in Fig. 3. The scale height error follows a diurnal cycle and minimizes at hour 13, around the same time neutral density RMSE reaches a minimum. Shown with experiment 1, these improvements are attributed to updating neutral temperatures, explained in Eq. 2. Updated temperatures impact vertical distributions of each species according to the species specific scale height. Above sub-orbital altitudes collisions of the neutral gas species become infrequent, and molecular diffusion, a rather slow process, controls the vertical distributions of neutral species. Thus, through the time update, helium is better specified and additionally explains the 14 hour lag time between scale height errors reaching a minimum and helium composition errors achieving a minimum. Updated scale heights have a stronger




Fig. 2: Global helium composition RMSE spanning the top three TIEGCM pressure levels, altitudes 360 km ~ 530 km.

and more immediate impact on neutral densities than updated neutral winds. Updated neutral winds are expected to affect the horizontal transport of species, however, RO data assimilation appears to have limited impact on improving transport effects.




Fig. 3: Global scale height RMSE spanning the top three TIEGCM pressure levels, altitudes 360 km ~ 530 km.

To illustrate a globally that neutral densities and helium composition is improving, Fig. 4 shows global prior density errors at the second to top pressure level, and Fig. 5 shows prior helium errors at the second to top pressure level over the winter pole. The first twelve hours are plotted, followed by every other hour. The initial errors are seen at the first hour, with red regions indicating overestimates and blue indicating underestimates. For neutral densities, the reduced error is seen globally and it reaches a minimum at UTC12 on day 175. Errors grow and reduce in the mentioned diurnal cycle, but remains smaller than the initial errors. For helium, compositions are generally overestimated over the south pole, and errors reduce. However the minimum error doesn't occur at UT 12 on day 175, but rather at UT 4 on day 176.

At the initial time step, ensemble neutral density is positively biased from the control for most of the globe, and negatively biased over the south pole. This may seem counter intuitive, since the the truth run has greater driver magnitudes than the mean driver of the experiment ensembles. Note that these plots are in TIEGCM's vertical coordinates of constant pressure level, not at constant altitudes. Greater magnitude drivers effectively expand the grid altitudes, and smaller magnitude drivers shrink the grid altitudes. Thus this, with the exponential drop off in neutral density with altitude, explains why the prior errors are positively biased. Assimilating COSMIC RO increases temperature, expanding TIEGCM's grid altitudes for a given pressure level and better specifying neutral densities.

For helium composition, a positive bias exists over the pole, opposite the negative bias in neutral density. When the atmosphere is cooler, the atmosphere collapses, allowing for greater concentration helium at lower altitudes. Thus the cooler experiment ensembles will have a greater composition of helium than the hotter truth run. As temperatures are increased, scale height is better specified and helium is vertically transported upward, reducing compositions and better matching true helium compositions.

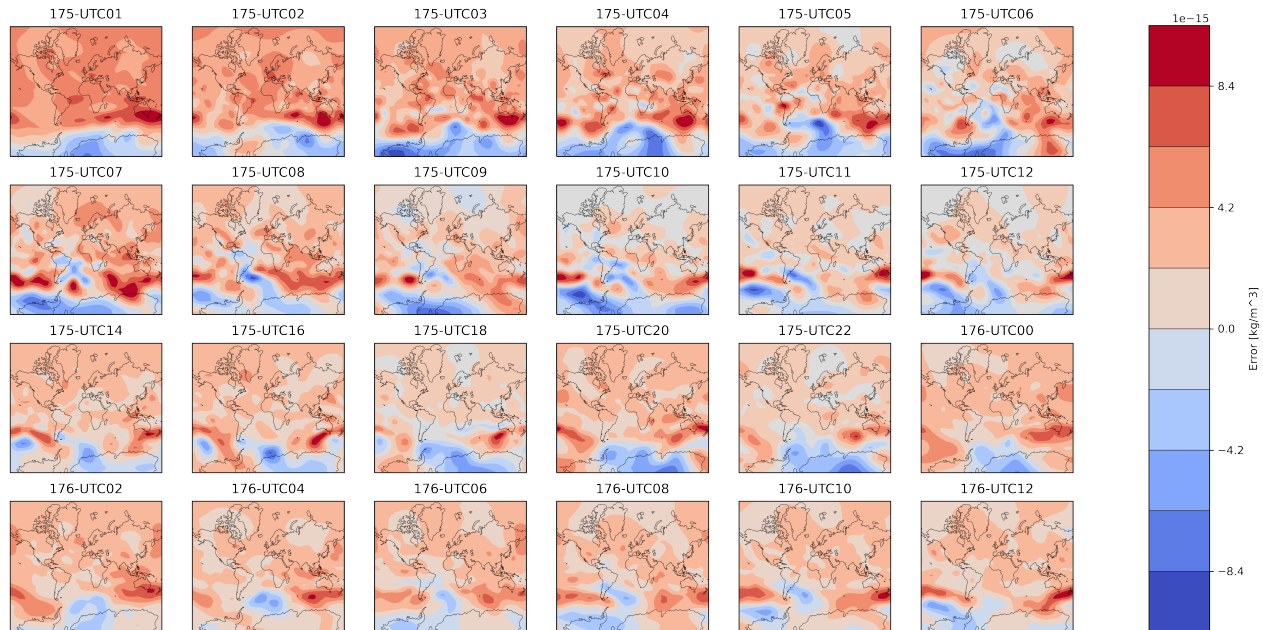


Fig. 4: Global neutral density errors for experiment 1 priors at the second to top pressure level of TIEGCM, altitudes 400 km ~ 530 km. Prior errors for each of first 12 hours is shown, followed by every other hour.

2.4 Discussion and Future Work

This study showcases the potential for COSMIC RO observations to estimate helium. COSMIC RO observations' ability to improve temperature estimates are captured in the EAKF through ensemble-based covariance information between electron densities and temperature. One manifestation of this strong coupling is how height of peak electron densities is affected by temperature. Temperature improvements in turn leads to the improvement in scale height estimates, enabling helium's vertical distributions to match the true values through an improved molecular diffusion processes in the time update. These updates improve neutral density globally, and is now shown to work well for altitudes that include many LEO satellites. These results obtained using simulated atmospheric data, and more work will be required to show this assimilation method is effective with *real* data. In this OSSE study, experiment results could directly be compared against known atmospheric states to evaluate the efficacy of the method. Validation of experiments using real data is challenging since there is little to no helium observation information, making it necessary to rely on limited derived neutral densities and orbit propagation errors. Future work will additionally involve running experiments during storm-time conditions when the atmosphere is more disturbed.

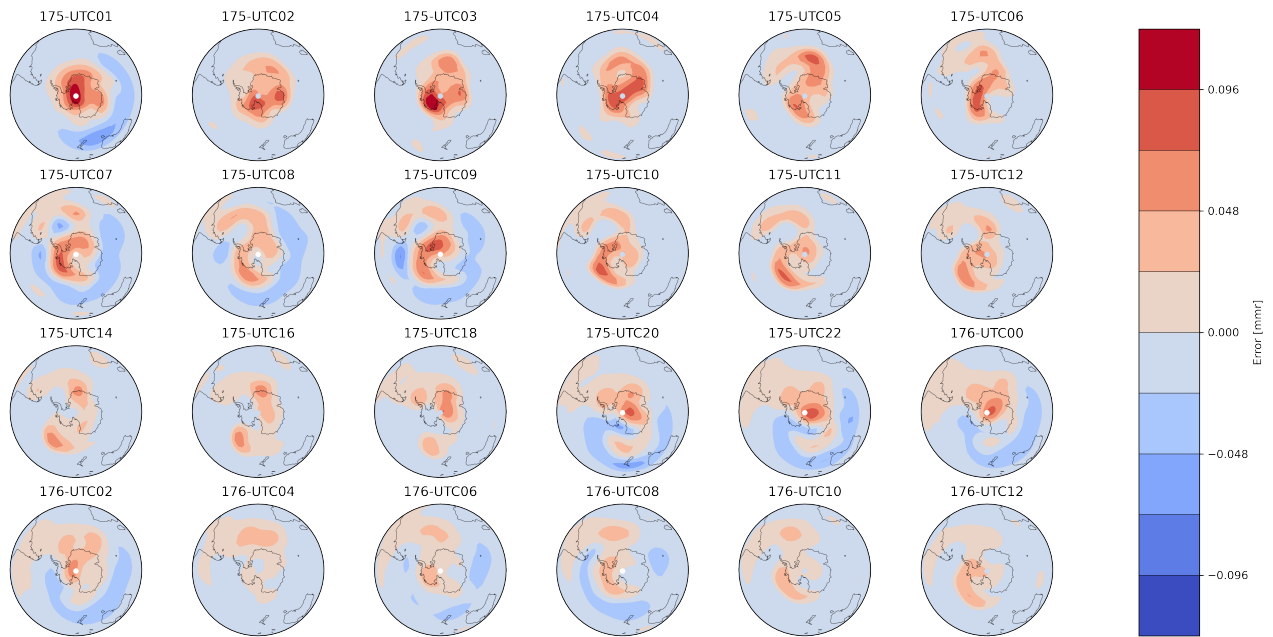


Fig. 5: Experiment 1 prior helium composition errors at the top layer of TIEGCM over the winter hemisphere pole, altitudes 400 km \sim 530 km. Prior errors for each of the first 12 hours is shown, followed by every other hour.

3. PARTICLE FILTER DRIVER ESTIMATION

While assimilating COSMIC RO observations improves neutral density estimation, it may not be enough to entirely correct neutral density errors. As shown in [4], orbit errors and neutral densities compared with CHAMP are improved, however, a model bias still exists with CHAMP observations. Driver estimation for TIEGCM is therefore needed to further improve model biases. The iterative reinitialization approach by [18] has shown a strong ability to eliminate TIEGCM model biases against CHAMP, but lacks an ability to estimate the driver distributions required for uncertainty quantification. To improve orbit position uncertainty quantification, estimating solar and magnetospheric drivers distributions are needed. The particle filter framework allows estimating the driver distribution to reduce model bias and quantify driver uncertainties.

To solve parameters estimation problems, Monte Carlo methods with Bayesian updates have been adapted to overcome the shortcomings of linear filters [8], [9], [6]. The particle filter is a widely used application of these methods, allowing the non-linear dynamics of the model and Bayesian updates to be used to estimate non-Gaussian distributions. Particle filters are also more commonly being adapted for geophysical systems [22]. A Gaussian distribution assumption is made for the EAKF, however, it is unknown how close this distribution matches the true driver distribution, and the uncertainties for atmospheric states. While a Gaussian assumption is easy to make, it is unclear how much this effects the distribution of atmospheric states and whether a better represented driver distribution would better represent atmospheric state uncertainties, potentially improving the efficacy of the EAKF and improving density uncertainties for orbit determination. Using a particle filter, a time-varying driver distribution is estimated and tested to see how well it can shift drivers to remove model biases in TIEGCM to match real observations. Once proven to be effective, this filter may be implemented to run online alongside the EAKF. For simplicity of the experiments, the two driver previously mentioned are estimated: F10.7 index and Kp index. Both of these indices have "true" observed values, and when used to drive TIEGCM give good magnitudes relative to true atmospheric states. However, the relationship of these parameters and atmospheric states are just proxies for the amounts of energy input into the upper atmosphere. Therefore, we view these parameters need to be calibrated. While inputting true indices into TIEGCM gives good priors for aligning with observation, estimated indices can still vary greatly from true indices as in providing the best representation of atmospheric states to observations.

Table 3: Particle Filter Variable Notation

Symbol	Meaning
F	TIEGCM Model
$\boldsymbol{\theta}^k(t_i)$	Driver states at time t_i , particle k
$\mathbf{f}^k(t_i)$	TIEGCM grid at time t_i , particle k
h	Neutral density and interpolation forward model
\mathbf{R}	Observation variance
$\mathbf{y}(t_i)$	Observation vector at time t_i
$\boldsymbol{\varepsilon}$	Driver process noise
$\mathbf{r}^k(t_{i+1})$	innovation vector at time t_{i+1} for particle k

3.1 Particle Filter Formulation

A bootstrap particle filter approach is adapted to implement the particle filter. The goal is to estimate drivers, however, observations are for neutral densities. As the relationship among F10.7 index and Kp index and neutral densities can be very complicated, depending on both drivers and the initial conditions of atmospheric states, TIEGCM is used in a unique way, acting as time-varying forward model for drivers. For each time window, TIEGCM will take in drivers, $\boldsymbol{\theta}(t_i)$, and initial atmospheric states, $\mathbf{f}(t_i)$, and propagate N particles forward from time t_i to time t_{i+1} . Time updated TIEGCM states through this propagation time window are compared to observations at the same hours, $\mathbf{y}(t_{i+1})$, and used to calculate the likelihood of each driver. To perform the measurement update, Bayes' theorem yields the posterior probability distribution relation

$$p(\boldsymbol{\theta}(t_i); \mathbf{f}(t_i) | \mathbf{y}(t_{i+1})) \propto p(\mathbf{y}(t_{i+1}) | \boldsymbol{\theta}(t_i); \mathbf{f}(t_i)) p(\boldsymbol{\theta}(t_i); \mathbf{f}(t_i)) \quad (9)$$

Particle weights may be calculated and normalized using sum of all weights, avoiding having to calculate the denominator in Bayes' theorem

$$w_i^k = \frac{p(\boldsymbol{\theta}^k(t_i); \mathbf{f}^k(t_i))}{\sum_{j=1}^N p(\boldsymbol{\theta}^j(t_i); \mathbf{f}^j(t_i))} \quad (10)$$

for all particles $k = 1, \dots, N$. Using a bootstrap method, weights are assigned to each particle and resampled based on these assigned weights. After resampling, particle weights are set to be equal, $1/N$ and used as the prior for the next measurement update.

It is noted that while the likelihood function depends on both the input drivers *and* the initial conditions, but only the estimated drivers are the goal of this estimation. From this, an important assumption is made: the posterior distribution for drivers is approximately independent of its initial conditions.

$$p(\boldsymbol{\theta}(t_i)) \approx p(\boldsymbol{\theta}(t_i); \mathbf{f}(t_i)) \quad (11)$$

While this approximation almost certainly is not true at the start time of the particle filter after many cycle run times the filter will be assumed to *forget* its initial conditions. This forgetting property is important to be assumed to make conclusions on the estimated driver distribution [6]. How valid this assumption is for TIEGCM will need to be further investigated, and can vary for each driver.

In this particle filter formulation, it is important to note that TIEGCM is being used differently than with EAKF described earlier. Previously, atmospheric states were being estimated and propagated with TIEGCM acting as a dynamical model. For the driver estimation formulation, the drivers have no inherent dynamical system. Observations are in the form of neutral densities, so to calculate the innovation vector the forward model must take in drivers and initial states, and output neutral densities. By running TIEGCM for some time span allows the drivers states to map to neutral densities using TIEGCM's non-linear dynamics. Thus, TIEGCM is acting as a time-dependent forward model, enabling driver states mappings to atmospheric states.

3.2 Particle Filter Algorithm

The implementation of the driver estimation particle filter is detailed below.

3.2.1 Initialization

To give ample spread among initial particle atmospheric states, each particle needs to be run through TIEGCM until they reach steady-state. This entails a 1-2 week long model run that serves as a spin-up period. This time period allows the drivers to fully impact the system and gives proper spread to atmospheric states. Like all filter initializations, an important design to consider is what prior distribution to use. While Gaussian fits were used for the EAKF, historical data suggests a highly skewed distribution. The year 2008 is considered an historic solar minimum, with extremely low F10.7 index values observed.

To best match this skewed distribution, a gamma distribution is fit to the yearly data of 2008 to provide the prior distribution for sampling. The gamma distribution was fit and adjusted to have the mean and standard deviations to match historical data. The skewed distribution provides many samples with low forcing conditions, while containing some high valued driver to capture storms should they occur. After sampling $N = 500$ sets of drivers, each particle is run with TIEGCM, F , to the start time of the particle filter.

$$\mathbf{f}^k(t_i) = F_{t \rightarrow t_i}(\boldsymbol{\theta}^k(t_0); \mathbf{f}^k(t_0)) \quad k = 1, \dots, N \quad (12)$$

3.2.2 Time Update

For the first cycle, TIEGCM is used to propagate each set of states k forward to the next update time by a time step Δt hours.

$$\mathbf{f}^k(t_{i+1}) = F_{t_i \rightarrow t_{i+1}}[\boldsymbol{\theta}^k(t_i); \mathbf{f}^k(t_i)] \quad k = 1, \dots, N \quad (13)$$

As was mentioned previously, TIEGCM is being used as a time-dependent forward model. While *atmospheric states* are updated in this step, updating drivers is done jointly in the time-update and measurement update steps. This step is by far the most computationally intensive step and requires the use of high-performance computing systems to realistically run.

3.2.3 Measurement Update

As the drivers take time to impact the system, observations at times following the onset of the filter cycle will be used to estimate particle weights. These observations used will include every time after but not including t_i , up to time t_{i+1} . These constitute m observations within the model run time that will be used to estimate the weights. The innovation vector used within the likelihood function is constructed for each particle, $\mathbf{r}^k(t_{i+1}) \in \mathbb{R}^m$.

$$\mathbf{r}^k(t_{i+1}) = \mathbf{y}(t_{i+1}) - h(F_{t_i \rightarrow t_{i+1}}[\boldsymbol{\theta}^k(t_i); \mathbf{f}^k(t_i)]) \quad (14)$$

The forward operator, h , here acts on the time propagated TIEGCM input with drivers $\boldsymbol{\theta}^k(t_i)^k$ and initial states $\mathbf{f}^k(t_i)$. The forward operator performs two functions: calculates the neutral densities at TIEGCM grid locations and interpolates them to observation locations. Neutral densities, ρ [kg/m³], are calculated using the ideal gas law

$$\rho = \frac{P}{RT \left(\frac{\gamma_O}{16} + \frac{\gamma_{O_2}}{32} + \frac{\gamma_{He}}{4} + \frac{1 - \gamma_O - \gamma_{O_2} - \gamma_{He}}{28} \right)} \quad (15)$$

where P is the pressure [N/m²], T is the temperature [K], $R = 8.314$ [J/K/kmol] is the universal gas constant, and $\bar{M} = \frac{\gamma_O}{16} + \frac{\gamma_{O_2}}{32} + \frac{\gamma_{He}}{4} + \frac{1 - \gamma_O - \gamma_{O_2} - \gamma_{He}}{28}$ [kg/kmol] is the mean molecular mass. γ_O , γ_{O_2} , and γ_{He} are the mixing ratio for O, O₂ and He, respectively. A simple vertical linear interpolation is done to find the neutral density at the observation altitude.

Once the innovation vectors $\mathbf{r}^k(t_{i+1})$ are calculated, they are fed into the likelihood function to compute particle weights. A Gaussian likelihood function is used for simplicity and for its ease in working with log-likelihoods.

$$w_i^k = p(\mathbf{y}(t_{i+1}) | \boldsymbol{\theta}^k(t_i); \mathbf{f}^k(t_i)) = \frac{1}{(2\pi)^{m/2} \det \mathbf{R}_i^{1/2}} \exp \left[-\frac{1}{2} (\mathbf{r}^k(t_{i+1}))^T \mathbf{R}_{i+1}^{-1} (\mathbf{r}^k(t_{i+1})) \right] \quad (16)$$

Weights are then normalized using Eqn. 10 and from the bootstrap method assumption setting the prior weights to $p(\boldsymbol{\theta}^k(t_i); \mathbf{f}^k(t_i)) = 1/N$, we can find the corresponding particle weights w_i^k and normalize them to \tilde{w}_i^k .

3.2.4 Resampling

After particle weights are assigned, particles are resampled based on their appropriate weights. Particles are then perturbed slightly using a Gaussian distribution to increase driver spread, $\varepsilon \sim N(0, \sigma^2)$.

$$\begin{bmatrix} \boldsymbol{\theta}_{F10.7}^k(t_{i+1}) \\ \boldsymbol{\theta}_{Kp}^k(t_{i+1}) \end{bmatrix}_{perturbed} = \begin{bmatrix} \boldsymbol{\theta}_{F10.7}^k(t_{i+1}) \\ \boldsymbol{\theta}_{Kp}^k(t_{i+1}) \end{bmatrix}_{resampled} + \begin{bmatrix} \varepsilon_{F10.7} \\ \varepsilon_{Kp} \end{bmatrix} \quad (17)$$

As mentioned previously, while the likelihood is meant to estimate the drivers, there is an entangling of the likelihood with the initial conditions. After resampling and perturbing, we now have a brand new set of N drivers, but the old set of initial TIEGCM states remain the same. For running the next time update using TIEGCM, these set of drivers need to be assigned to our previous particle initial conditions. This is done by formulating this as an optimization assignment problem optimization to deterministically align both sets of drivers. Generally, higher valued resampled drivers will need to be aligned with higher valued prior drivers, and lower valued resampled drivers aligned with lower valued prior drivers. A cost matrix, \mathbf{C} , between the old list of drivers and the new list of resampled drivers is generated using the normalized Euclidean norm. The optimization problem is constructed as follows.

$$\min_S \sum_{i=1}^u \sum_{j=1}^v C_{ij} S_{ij} \quad (18)$$

$$\text{s.t.} \quad \sum_{i=1}^v S_{ij} = 1, \quad j = 1, \dots, u \quad (19)$$

$$\sum_{j=1}^u S_{ij} = 1, \quad j = 1, \dots, v \quad (20)$$

$$S_{ij} \in \{0, 1\} \quad (21)$$

The problem is solved using the Hungarian algorithm. With a new set of drivers ordered to their proper TIEGCM states, the filter moves to the time update for the next time window and repeats.

3.3 Algorithm Outline

1. Initialize N particles with prior driver distribution for a spin-up period of 15 days to the initial start time at t_i .
2. Propagate each particle through TIEGCM from t_i to t_{i+1} using inputs $\boldsymbol{\theta}^k(t_i)$ and $\mathbf{f}^k(t_i)$. Save output at the frequency observations occur.
3. Collect observations after but not including t_i to t_{i+1} . Find TIEGCM states at the corresponding observation location and time for each particle to produce $\mathbf{r}^k(t_i)$.
4. Evaluate the likelihood function for each particle to calculate its weight w_i^k , computed in log-likelihood space.
5. Normalize each weight in log space, and convert from log space, \tilde{w}_i^k .
6. Resample N particles $\boldsymbol{\theta}^k(t_{i+1})$ according to their weights \tilde{w}_i^k , and perturb using $\varepsilon \sim N(0, \sigma^2)$.
7. Assign driver particles $\boldsymbol{\theta}^k(t_{i+1})$ to TIEGCM initial states $\mathbf{f}^k(t_{i+1})$ and move to the next time window. Repeat process and go back to step 2.

3.4 Particle Filter Degeneracy

Like all particle filters, the major filter issue to address is filter degeneracy. While dimensional size of the state space is small at two drivers, the size of the observation space is very large for a particle filter. In the observation space for neutral density observations, there are commonly around 77 CHAMP observations per hour and 12 hours of data will be used each cycle. Each of these observations are considered independent, so multiplying all these likelihoods together quickly cause the likelihood to vanish beyond machine precision. While log-likelihood space minimizes machine precision errors, using that many observations quickly causes weights to collapse. As [16] explains about high-dimensional particle filters, the filter overestimates the information the observations provide and finds the

intersection of all the likelihoods, causing the weight to quickly collapse to a small number of particles. There were four methods used to avoid particle filter degeneracy: 1) Compute weights in the log-likelihood space, 2) increase the particle number as feasibly high, 3) dramatically increase observation variance and 4) introduce driver perturbation.

Due to the very small likelihoods produced from over 900 independent observations, a Gaussian likelihood assumption nicely allows the log-likelihoods to be computed and each observation likelihood added together. Each weight is calculated in the log space and added together.

$$\log w_i^k = \sum_{j=1}^m -\frac{1}{2} \left[\log 2\pi + \log \sigma_j + \left(\frac{r_j^k(t_i)}{\sigma_j} \right)^2 \right] \quad (22)$$

The weights are then normalized in log space via

$$\log \tilde{w}_i^k = \log w_i^k - \log \sum_{j=1}^N w_i^j \quad (23)$$

Where the total addition in log space may be done by shifting all the log weights by the largest value (all values are negative). Once the new normalized weights are calculated, each particle weight is converted back from log space. These log operations are very convenient to perform since the specific weights produced by the likelihoods do not matter as much as a particle's weight relative to all other particles.

The next proposition is to increase the number of particles as much as possible. The number of particles is feasibly limited to the order of hundreds of particles due to limits posed by TIEGCM. It important to iterate that TIEGCM is a very computationally intensive model, containing over half a million states, and is being run hundreds of times each filter cycle. Running TIEGCM requires a high-performance computing to solve the model. 500 particles are run within the filter which should provide ample spread for two drivers, however may still not be enough for providing ample state spread.

The most impactful proposition in avoiding filter degeneracy is to artificially inflate the observation uncertainties for calculating likelihoods. Aside from increasing particle size, the size of the observation uncertainties are the biggest influence in preventing filter degeneracy. However, this approach is also the most ad hoc. What values to choose to model these uncertainties is not yet understood. Not only does the uncertainty need to account for observation variance in estimating the *true* atmosphere, it more importantly needs to relate how useful that observation is to estimating TIEGCM drivers. Correlations between among neutral densities and TIEGCM drivers will need to be much further investigated to understand their relationship. These correlations differ along latitudes and local times, and the chosen uncertainty must be representative for all drivers. As an ad hoc solution, uncertainties are found by multiplying the true observation uncertainty by a factor of 10^4 in response to using over 900 observations per filter cycle. This value works well to ensure particle weights do not collapse while not fully drowning out observation information.

Lastly, driver perturbations are added not only to avoid degeneracy, but to allow the driver distribution to evolve over time and follow true observed changing drivers. If no driver noise were added, weights would almost assuredly converge a few particles. While the determined process noise added will depend on real space weather variability, adding some relatively small perturbations to sample drivers will allow particle filter experiments to converge to drivers not represented in the prior distributions. The frequency at which noise is added will need to be further tuned to avoid degeneracy, but also prevent driver distributions from devolving into noise.

3.5 Experiment Results

The particle filter is initialized with a prior based on historical data for 2008. Each particle is run for a spin-up period of 15 days to the starting time at UT 0 on June 23rd, 2008. The filter is run for four days, with drivers updated every 12 hours, and using observations from all 12 hours. During the filter run time, there are initially low Kp indices, with a minor storm occurring on June 26th, and historically low F10.7 indices. Gaussian noise was used to perturb drivers after resampling with $\epsilon_{F10.7} \sim N(0, 2^2)$ and $\epsilon_{Kp} \sim N(0, 0.2^2)$. Neutral densities estimated in [19] are used as true observations. The initial variances that were calculated with this data set are additionally used, and multiplied by a factor of 10^4 to account for the uncertainties between neutral density states and the solar drivers. The driver weighted means are calculated using Eq. 24. Comparisons between true drivers and estimated drivers are seen in Fig. 6.

$$\mu_\theta = \frac{\sum_{i=1}^N \tilde{w}_i \theta_i}{\sum_{i=1}^N \tilde{w}_i} \quad (24)$$




Fig. 6: Observed Kp and F10.7 indices over the filter run period compared with the the particle filter estimated Kp and F10.7 indices.

The prior particle distribution and each subsequently estimated particle distribution is seen in Fig. 7. The prior begins with considerable spread, and particle weights quickly collapses to small Kp values, and moves towards lower F10.7 values. The distribution looks to become more noisy and more Gaussian shaped.

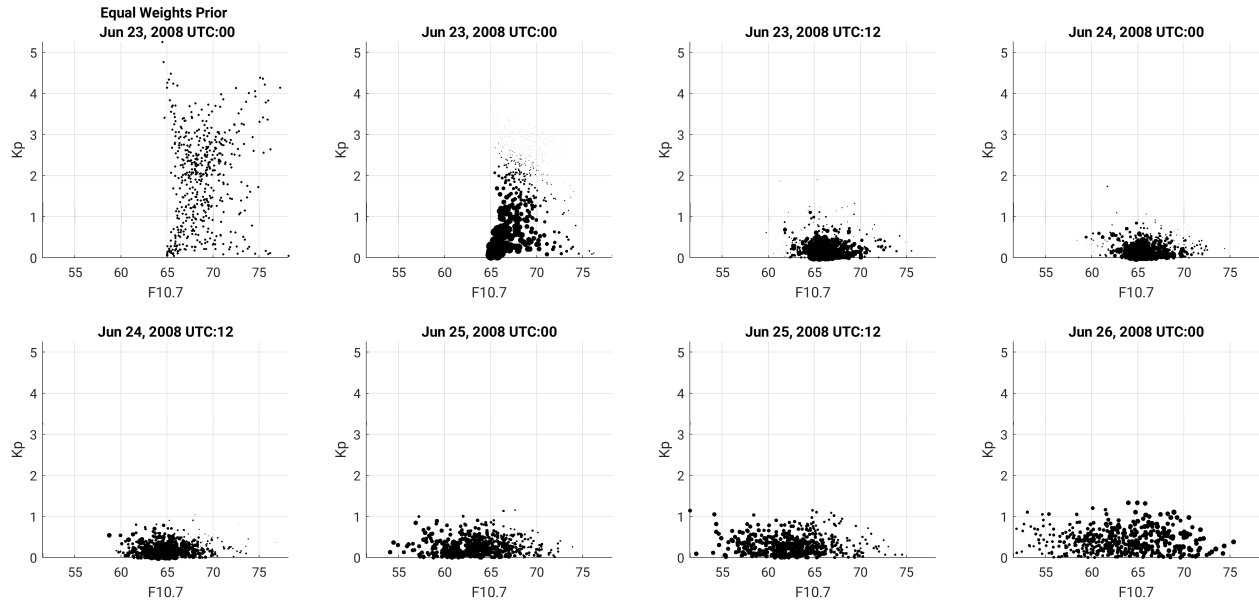


Fig. 7: 2D Particle driver distributions over the particle filter run time, updated every 12 hours. The first plot is the equally weighted prior sampled from a gamma distribution. Particle size is relative to the weight of that particle.

After a few filter cycles, the particle filter improves neutral density estimates as shown in Fig. 8. How the weighted mean is evolving over time in comparison with CHAMP data may be seen in Fig. 9. CHAMP's orbit is polar at an altitude around 340 km. Following the +90 degree latitude, CHAMP is in a descending orbit on the night-side, and following the -90 degree latitude, CHAMP is in an ascending orbit on the day-side. Peak neutral densities are observed at low latitudes just above the day-side equator, and minimum densities are observed at the winter pole. Even with a good prior based on historical data, the weighted mean is still biased above an historically low neutral density values. This bias illustrates how TIEGCM's drivers are acting as proxies, and inputting real driver values may still produce model biases. By the third day, filter's estimated mean lowers to meet slightly rising CHAMP densities, providing good agreement on June 25th. Neutral density spreads are seen at a few time locations through boxplots. Distributions for June 24th show a positively skewed distribution, as reflected in the skewed prior driver distribution.

Neutral densities boxplots collapse more with more filter time as the driver distribution collapses more around the optimal drivers. Boxplots over peak neutral densities at low latitudes on the day-side show greater spread, while spreads at minimum densities over the pole show the least spread. These reflect where drivers are most impactful in specifying neutral density and can be calibrated for increasing density spread.

While the particle filter matches well with CHAMP neutral densities on the third day of the filter, observations begin to diverge following UT 19 on June 25th. Additionally shaded regions for Fig. 9 show the times when the Kp index is greater than 3. Following the onset of higher geomagnetic activity at UT18 on June 25th, neutral density observations rise on the day-side northern hemisphere while F10.7 index remains extremely low. The filter still matches well with night-side observations, but fails to extend to capture the peak densities. With estimated Kp indices largely collapsing to less than 1, there are not as many particles with higher Kp values to capture these changes. While there is still some day-side density spread encompassing these values, the particles are not capturing densities as well as on the night-side. This is the trade-off of using 12 hours of data to calculate particle weights, as the measurement update frequency may not be high enough to capture smaller time-scale changes in neutral density.




Fig. 8: TIEGCM weighted mean RMSE against CHAMP neutral densities.

The estimated distribution is difficult in assessing whether or not the distribution is converging to Gaussian can be assessed. Adding noise after resampling is likely to be moving the driver distribution to towards a Gaussian distribution. Additionally, while inflating \mathbf{R} is helping to prevent the weights from collapsing, over-inflating will cause too much information from TIEGCM to be lost to noise. This will happen if TIEGCM's response to new drivers isn't faster than the rate at which noise is added. This is likely true with resampling occurring every 12 hours. To assess the estimated driver distribution's resemblance to a Gaussian one, Q-Q plots were generated for both F10.7 and Kp in Fig. 10. The initial prior is gamma, and this is seen in the Q-Q plots a right skewed distribution. For F10.7, we see samples increasingly align with the black line over each cycle, indicating the distribution is fitting to a Gaussian distribution. However, for Kp, the distribution begins right skewed, and remains right skewed, retaining its non-Gaussian structure even with added noise. These findings suggest geomagnetic drivers have non-linear effects on neutral densities, while the EUV driver for F10.7 has more of a linear effect on neutral densities. This non-linearity could explain why the filter had a greater difficulty capturing changes in Kp, and was better at converging on F10.7 values.

3.6 Discussion and Future Work

Overall, this experiment has proved the particle filter's capability to improve neutral densities through estimating drivers. While previous work using the iterative reinitialization approach from [18] has also shown an ability to reduce model biases by improving driver estimates, this work seeks to extend driver estimation capabilities. The particle filter framework allows for a more formal approach that is amenable to the estimation of non-Gaussian driver distributions. Using a particle filter framework to estimate drivers provides the capability to estimate driver distributions and perform uncertainty quantification. There are many parameters that need to be better calibrated to improve filter performance: resampling frequency, observation variance and driver perturbations. Changing these values could drastically affect particle weights and how quickly the driver distributions may devolve into noise or collapse to a single particle.

Many of these challenges stem from relying on TIEGCM as a time-dependent forward model. TIEGCM is a complex, high-dimensional model that requires considerable computationally resources, so particle numbers are severely limited by the CPU, memory and storage requirements of TIEGCM simulations. Additionally state response times in TIEGCM with new drivers is a highly non-linear relationship and benefits from further studied using different re-

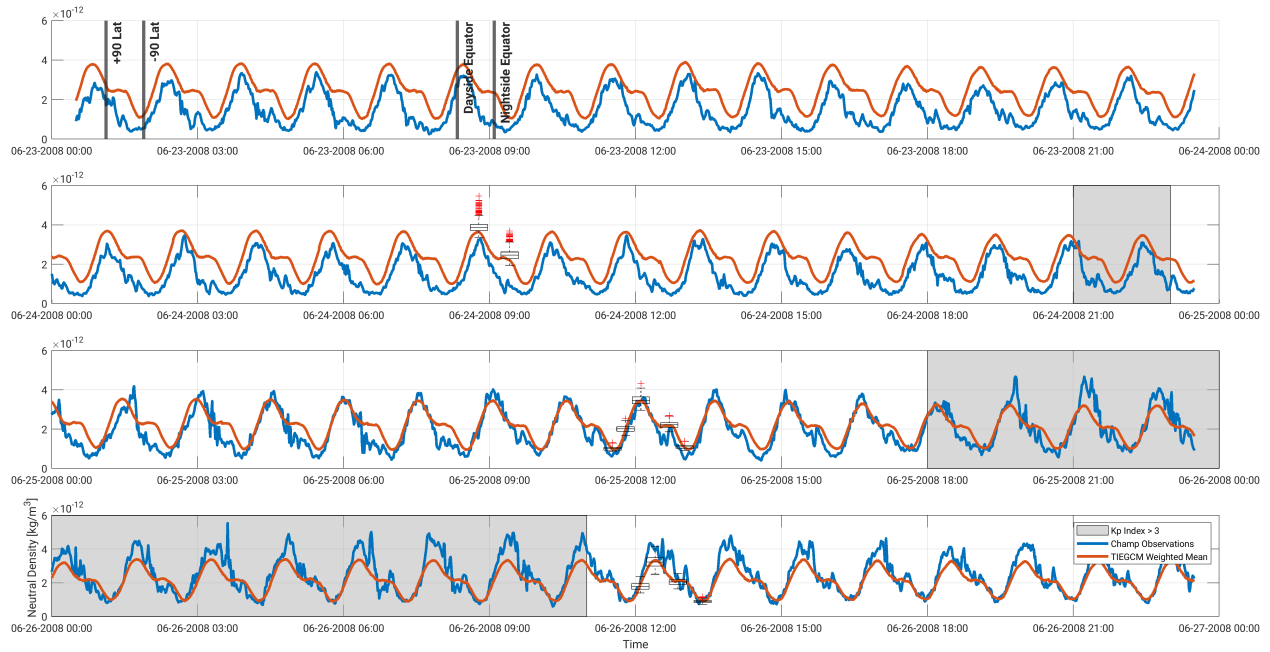


Fig. 9: CHAMP neutral density observations and their corresponding weighted mean TIEGCM locations. Drivers are updated every 12 hours, with the first resampling at UT12 on June 23rd. CHAMP is in a descending orbit on the night-side following +90 degree latitude and in an ascending orbit on the day-side following the -90 degree latitude and has orbit period of 90 minutes. Box plots are unweighted particle spread at the corresponding time. Shaded regions are periods when there is higher geomagnetic forcing, $K_p > 3$.

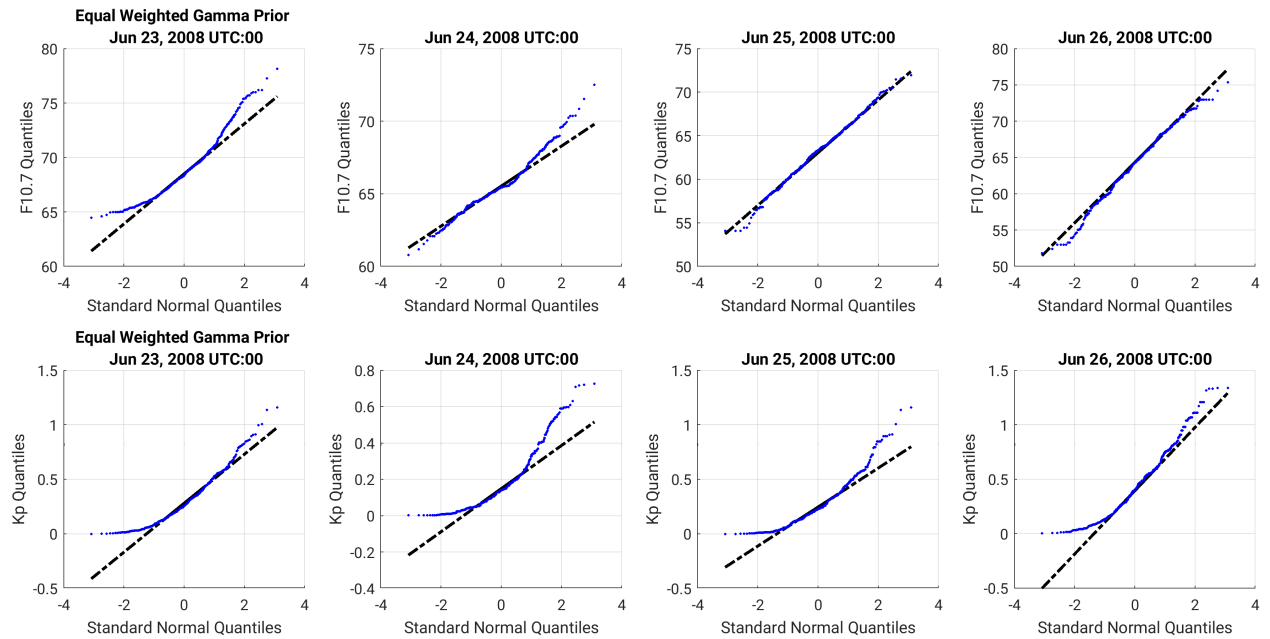


Fig. 10: Q-Q plot compared with the standard normal distribution. Top row is for F10.7 and bottom row is for K_p . The leftmost plots are the prior Gamma distributions, and each subsequent horizontal plot is the estimated distribution every 24 hours.

sampling frequencies. There is a delicate balance between capturing fast-time scale events, and keeping the driver distributions from degenerating into noise. Even with a good prior, an appropriate balance is difficult to achieve.

Some potential solutions may be an adaptive inflation, that will vary driver noise based on the degree of bias in certain regions of the globe, and allow for region specific weights associated with each particle. These features will be helpful to capture the thermosphere during storm-times, when rapid changes can occur through solar flare or geomagnetic storms.

This particle filter is developed with the goal of computing end to end uncertainty quantification for orbit determination. It is not clear what distribution shape drivers should take. Drivers are commonly assumed to be Gaussian, but this is a simple assumption that, at least is shown for Kp and F107 distributions, does not hold to be true. To better capture storm-time heating effects on the thermosphere, non-Gaussian driver distributions are needed. Perhaps magnetospheric drivers need to be represented more complex models, such as AMGeO [12], than Kp index. AMGeO estimates the 2-dimensional distribution of auroral electron precipitation and electrostatic potential from actual observations that may be used to specify geomagnetic forcing over TIEGCM's polar cap.

4. CONCLUSION

The two approaches presented in this paper provide a means of performing end-to-end uncertainty quantification for orbit errors. The first approach is assimilating COSMIC RO electron density observations into TIEGCM to update neutral states using strongly coupled ionosphere-thermosphere ensemble data assimilation. Experiments show improvements in neutral density and helium composition estimations due to better specified temperature distribution that is reflected in improvements in scale height estimations. The second approach looks to estimate solar and magnetospheric drivers using real in-situ density observations and TIEGCM as a time-dependent forward model. Initial results show the ability to reduce model bias, however, more work is needed in verifying the estimated driver distributions. Both the RO ensemble data assimilation and the particle filter driver estimation can be used to globally specify neutral densities, and facilitates a probabilistic representation of neutral density uncertainties during solar flares and geomagnetic storms so that LEO space traffic may be better managed.

5. ACKNOWLEDGEMENTS

This work is supported by the NSF award AGS-1848544. The DART software was obtained from NCAR (<https://dart.ucar.edu/>). The TIEGCM software was obtained from High Altitude Observatory, NCAR (<https://www.hao.ucar.edu/modeling/tgcm/tie.php>). NCAR is sponsored by the NSF. The COSMIC data were obtained from COSMIC Data Analysis and Archival Center (<https://www.cosmic.ucar.edu/what-we-do/cosmic-1/data/>). We acknowledge NCAR and UCAR for their community software and data support and the high-performance computing support from Cheyenne (doi:10.5065/D6RX99HX) provided by NCAR. We thank Dr. Eric Sutton for his assistance with CHAMP observations.

6. REFERENCES

- [1] J. L. Anderson. An ensemble adjustment Kalman filter for data assimilation. *Monthly Weather Review*, 129(12):2884–2903, 2001.
- [2] J. L. Anderson. A local least squares framework for ensemble filtering. *Monthly Weather Review*, 131(4):634–642, 2003.
- [3] T. E. Berger, M. J. Holzinger, E. K. Sutton, and J P Thayer. Flying through uncertainty. *Space Weather*, 18(1):e2019SW002373, 2020.
- [4] N. Dietrich, T. Matsuo, and C.T. Hsu. Thermospheric Neutral Density Specification and Forecasting via Driver Estimation and Assimilation of COSMIC Radio Occultation Data. AGU Fall Meeting, 2020.
- [5] A. E. Hedin, C. A. Reber, G. P. Newton, N. W. Spencer, H. C. Brinton, H. G. Mayr, and W. E. Potter. A global thermospheric model based on mass spectrometer and incoherent scatter data MSIS, 2. Composition. *Journal of Geophysical Research*, 82(16):2148–2156, 1977.
- [6] N. Kantas, A. Doucet, S. S. Singh, J. Maciejowski, and N. Chopin. On particle methods for parameter estimation in state-space models. *Statistical Science*, 30(3):328–351, 2015.
- [7] G M Keating and E J Prior. The winter helium bulge. 1968.
- [8] G. A. Kivman. Sequential parameter estimation for stochastic systems. *Nonlinear Processes in Geophysics*, 10(3):253–259, 2003.

- [9] A. Laine, M. Solonen and H. Haario, H. and Järvinen. Ensemble prediction and parameter estimation system: The concept. *Quarterly Journal of the Royal Meteorological Society*, 138(663):289–297, 2012.
- [10] H. Liu, H. Lu, S. Watanabe, W. Ko, and C. Manoj. Contrasting behavior of the thermosphere and ionosphere in response to the 28 October 2003 solar flare. 112(A7), 2007.
- [11] T. Matsuo and C.H. Hsu. Inference of Hidden States by Coupled Thermosphere-Ionosphere Data Assimilation: Applications to Observability and Predictability of Neutral Mass Density. *Upper Atmosphere Dynamics and Energetics*, pages 343–363, 2021.
- [12] T. Matsuo, L.M. Kilcommons, W. Mirkovich, J.M. Ruohoniemi, S. Chakraborty, B.J. Anderson, and S.K. Vines. Assimilative Mapping of Geospace Observations (AMGeO): Unified Global and Local Perspectives on High-latitude Ionospheric Electrodynamics. In *AGU Fall Meeting Abstracts*, volume 2020, pages SA026–03, dec 2020.
- [13] L. Qian, A. G. Burns, B. A. Emery, B. Foster, G. Lu, A. Maute, A. D. Richmond, R. G. Roble, S. C. Solomon, and W. Wang. The NCAR TIE-GCM: A community model of the coupled thermosphere/ionosphere system. *Modeling the ionosphere-thermosphere system*, 201:73–83, 2014.
- [14] C. A. Reber, D. N. Harpold, R. Horowitz, and A. E. Hedin. Horizontal distribution of helium in the earth’s upper atmosphere. *Journal of Geophysical Research*, 76(7):1845–1848, 1971.
- [15] A. D. Richmond, E. C. Ridley, and R. G. Roble. A thermosphere/ionosphere general circulation model with coupled electrodynamics. *Geophysical Research Letters*, 19(6):601–604, 1992.
- [16] C. Snyder, T. Bengtsson, P. Bickel, and J. Anderson. Obstacles to high-dimensional particle filtering. *Monthly Weather Review*, 136(12):4629–4640, 2008.
- [17] Mark F. Storz, Bruce R. Bowman, James I. Branson, Stephen J. Casali, and W. Kent Tobiska. High accuracy satellite drag model (HASDM). *Advances in Space Research*, 36(12):2497–2505, 2005.
- [18] E. K. Sutton. A New Method of Physics-Based Data Assimilation for the Quiet and Disturbed Thermosphere. *Space Weather*, 16(6):736–753, 2018.
- [19] E. K. Sutton, R. S. Nerem, and J. M. Forbes. Density and Winds in the Thermosphere Deduced from Accelerometer Data. 44(6), 2007.
- [20] E. K. Sutton, J. P. Thayer, W. Wang, S. C. Solomon, X. Liu, and B. T. Foster. A self-consistent model of helium in the thermosphere. *Journal of Geophysical Research A: Space Physics*, 120(8):6884–6900, 2015.
- [21] J. P. Thayer, X. Liu, J. Lei, M. Pilinski, and A. G. Burns. The impact of helium on thermosphere mass density response to geomagnetic activity during the recent solar minimum. *Journal of Geophysical Research: Space Physics*, 117(7):1–14, 2012.
- [22] P. J. Van Leeuwen. Particle filtering in geophysical systems. *Monthly Weather Review*, 137(12):4089–4114, 2009.