Space Weather®

Key Points:

- Assimilative Canadian High Arctic Ionospheric Model (A-CHAIM) is an operational near-real-time data assimilation providing an improved ionospheric electron density at high latitudes
- The model shows improved performance in the topside ionosphere, as well as in F2-layer peak parameters (foF2, hmF2)
- A-CHAIM is the first operational use of a particle filter for space environment specification

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

B. Reid,
breid.phys@gmail.com

A-CHAIM: Near-Real-Time Data Assimilation of the High Latitude Ionosphere With a Particle Filter

Ben Reid[1] ![ORCID], David R. Themens[1,2] ![ORCID], Anthony McCaffrey[1] ![ORCID], P. T. Jayachandran[1] ![ORCID], Magnar G. Johnsen[3], and Thomas Ulich[4] ![ORCID]

[1]Department of Physics, University of New Brunswick, Fredericton, NB, Canada, [2]School of Engineering, University of Birmingham, Birmingham, UK, [3]Tromsø Geophysical Observatory, UiT the Arctic University of Norway, Tromsø, Norway, [4]Sodankylä Geophysical Observatory, Sodankylä, Finland

Abstract The Assimilative Canadian High Arctic Ionospheric Model (A-CHAIM) is an operational ionospheric data assimilation model that provides a 3D representation of the high latitude ionosphere in Near-Real-Time (NRT). A-CHAIM uses low-latency observations of slant Total Electron Content (sTEC) from ground-based Global Navigation Satellite System (GNSS) receivers, ionosondes, and vertical TEC from the JASON-3 altimeter satellite to produce an updated electron density model above 45° geomagnetic latitude. A-CHAIM is the first operational use of a particle filter data assimilation for space environment modeling, to account for the nonlinear nature of sTEC observations. The large number ($>10^4$) of simultaneous observations creates significant problems with particle weight degeneracy, which is addressed by combining measurements to form new composite observables. The performance of A-CHAIM is assessed by comparing the model outputs to unassimilated ionosonde observations, as well as to in-situ electron density observations from the SWARM and DMSP satellites. During moderately disturbed conditions from 21 September 2021 through 29 September 2021, A-CHAIM demonstrates a 40%–50% reduction in error relative to the background model in the F2-layer critical frequency (foF2) at midlatitude and auroral reference stations, and little change at higher latitudes. The height of the F2-layer (hmF2) shows a small 5%–15% improvement at all latitudes. In the topside, A-CHAIM demonstrates a 15%–20% reduction in error for the Swarm satellites, and a 23%–28% reduction in error for the DMSP satellites. The reduction in error is distributed evenly over the assimilation region, including in data-sparse regions.

Plain Language Summary While we often think of space as a perfect vacuum, the region of space near Earth is filled with plasma, known as the ionosphere. This plasma can have significant effects on satellites and radio communications, and so it is important to be able to detect changes in the ionosphere. The Assimilative Canadian High Arctic Ionospheric Model (A-CHAIM) is a new system that has been developed to help improve our understanding of space weather in the northern hemisphere. It combines data from several different kinds of instrument to produce a forecast to predict the local space environment for the next 2 hours. One of the most important data sources used in A-CHAIM is Global Positioning System (GPS) stations. Changes to the ionosphere disrupt GPS service, but we can use these disruptions to learn how the plasma is moving. These observations require special processing to be useful, and so new techniques had to be developed for A-CHAIM. We compare the predictions made by A-CHAIM to measurements of the space plasma from satellites, and specialized instruments that use radio signals to measure the ionosphere from the ground. This allows us to show that A-CHAIM is able to produce an improved space weather forecast.

1. Introduction

The high latitude ionosphere has historically been a challenging system to model (Buchert, 2020; Lockwood et al., 1990; Rasmussen et al., 1986). A rich collection of external drivers and interactions drive ionospheric behavior, including strong electric fields, magnetospheric coupling via particle precipitation and current systems, and rapid changes in the thermospheric state. These dynamic conditions, paired with a lack of high latitude observations when compared to mid and low latitudes, present a substantial problem for operational ionospheric modeling. With increased interest in polar ionospheric monitoring (Thayaparan et al., 2018) and High Frequency (HF) communications, it is now imperative that a near-real-time, operational model of high latitude electron density be developed and deployed for use in this region.

Conventional physics-based models have generally struggled to perform independently at sufficient accuracies in their specification of electron density for operational applications, when compared to empirical models (Shim et al., 2018). This is due in part to limited spatial resolution of prescribed electric fields and particle precipitation (Cosgrove & Codrescu, 2009) and the general quality of model driver specification (Fernandez-Gomez et al., 2019). Even at mid and low latitudes these challenges, and the computational requirements of physics-based models, have often led to operational users having to rely on empirical ionospheric models, such as the International Reference Ionosphere (IRI) (Cervera et al., 2018; Cervera & Harris, 2014); and NeQuick (Montenbruck & González Rodríguez, 2019), which have been demonstrated to generally outperform most other available models (Shim et al., 2011, 2018). At high latitudes, however, the IRI insufficiently represents climatological behavior and nearly completely lacks specification of ionospheric storm-time variability (Bjoland et al., 2016; Themens & Jayachandran, 2016; Themens et al., 2014). This was highlighted in Themens et al. (2014) which showed that the IRI can exhibit errors in peak ionospheric critical frequency (foF2) in excess of 70% at times; furthermore, in Themens et al. (2020) the IRI was demonstrated to represent less than 5% of the amplitude and between 0.5% and 9% of the variance of ionospheric variability on intermediate (1-to-30 days) timescales.

These limitations inspired the development of the Empirical Canadian High Arctic Ionospheric Model (E-CHAIM), which was designed explicitly to better represent the climatological ionosphere at high latitudes (Themens et al., 2017; Themens, Jayachandran, & McCaffrey, 2019; Themens, Jayachandran, & Varney, 2018). The model generally exhibits strong performance in the polar cap, auroral zone, and Russian sector (Maltseva & Nikitenko, 2021; Themens, Jayachandran, McCaffrey, Reid, & Varney, 2019; Themens et al., 2021); however, it struggles at sub-auroral latitudes in the North American sector (Themens et al., 2021) and, despite doing better than the IRI at high latitudes, it is still only capable of representing up to 50% of the amplitude and 4%–25% of the variance of ionospheric variability on intermediate timescales (Themens et al., 2020). This ultimately necessitates the use of data assimilation to improve further upon E-CHAIM's representation over North America and to capture smaller spatial and temporal scales. The focus of this work is to develop a data assimilation technique that can be used as a near-real-time operational system to produce a higher fidelity 4D electron density model of the high latitude ionosphere, the Assimilative Canadian High Arctic Ionospheric Model (A-CHAIM).

There are many data assimilation techniques used in ionospheric research, and by the broader geophysical community, each with their own advantages and disadvantages (Angling et al., 2009; Bust & Immel, 2020; Chartier et al., 2016; Elvidge, Sean & Angling, Matthew J., 2019; Lee et al., 2012; Nickisch et al., 2016; Prol et al., 2021; Scherliess et al., 2004; Schunk et al., 2016; Spencer & Mitchell, 2003). Any choice of assimilation technique for A-CHAIM must be able to operate in near-real-time on reasonable hardware. This requirement to run in near-real-time places sharp constraints on what level of model complexity we are able to use. Running an ensemble of physics-based models, such as a General Circulation Model (GCM), would require significant computing resources. Existing physics-based assimilations at high latitudes have not shown strong improvements over climatology (Shim et al., 2011, 2018), and so it is difficult to justify using such a computationally expensive model if one is interested purely in electron density. Of course, physics-based data assimilation does have its advantages, as it can be used to infer information about other elements of the state space, such as thermospheric winds and densities (Chartier et al., 2016; Elvidge, Sean & Angling, Matthew J., 2019; Lee et al., 2012).

There exist a dozen or more different data assimilation models of the ionosphere that mainly use GNSS slant TEC measurements. Unfortunately, the path-integrated nature of these measurements heavily restricts the constraint afforded by them, particularly in the vertical distribution of electron density, and instrumental biases pose a substantial challenge, particularly at high latitudes (Coster et al., 2013; Nesterov & Kunitsyn, 2011; Prol et al., 2021; Themens et al., 2015).

The reconstruction of a 2D or 3D density using line integrals (tomography) is a common nonlinear inverse problem across many fields in medicine and geophysics. Tomographic techniques with radiofrequency beacon satellites for ionospheric studies have been practiced for decades (Prol et al., 2021). However, tomography works best when you have dense, evenly spaced networks of receivers (Chartier et al., 2014), which are not available in the high latitude region. Successful use of tomographic techniques usually requires careful conditioning and regularization, or fitting the solution to horizontal and vertical basis functions to reduce the dimensionality of the problem (Bust et al., 2004; Spencer & Mitchell, 2003).

There has been considerable success in using GNSS measurements to produce maps of vertical TEC (vTEC) as an operational product (Hernández-Pajares et al., 2009; Wielgosz et al., 2021). These products convert the

fundamentally nonlinear sTEC measurements into vTEC through projection functions, to bypass the limitations in the reconstruction of the vertical structure of electron density from GNSS measurements. While these products do have exceptional value for many user segments, they do not produce a 3D electron density model and they have limitations in their performance based on the projection required in their construction (Smith et al., 2008). Furthermore, these vTEC maps only use a single data type (GNSS measurements) in their reconstruction and are thereby unable to take advantage of complementary measurements.

To meet the computational limitations placed on a near-real-time system, we will here instead pursue the development of an empirical model-based data assimilation. A-CHAIM will be constructed as a series of perturbations on E-CHAIM. This sort of scheme has been used in an operational assimilation system before with the IRI Real-Time Assimilative Mapping (IRTAM), in which the diurnal and spatial profile parameters that govern the behavior of the IRI are adjusted using autoscaled ionosonde data (Galkin et al., 2012). IRTAM is unique in its approach of updating the coefficients of the background model, rather than using a conventional grid or voxel representation of its state space. This approach places some limitations on how the IRTAM can operate. The IRI basis set for its ionospheric peak parameters is made up of Fourier components in modified dip latitude, local time, and longitude (Jones & Gallet, 1962). Because local time is part of the horizontal basis set of the model, IRTAM requires a 24-hr time history of geographically-fixed data to be assimilated and cannot mix data types (Galkin et al., 2012). This is problematic, as the distribution of ionosondes is severely limited at high latitudes, and this approach precludes using more widely distributed GNSS data. Pignalberi et al. (2021) show that while IRTAM improves the representation of foF2 in the regions covered by data, it degrades performance in regions away from the ionosondes. Its reliance on ionosonde data also results in limited performance in hmF2, not just in regions away from data, but also in their vicinity. As such the IRTAM approach, while having desirable elements in terms of its computational efficiency and straightforward integration into the IRI, is not suitable for our application.

A-CHAIM uses a particle filter with 1,000 particles to assimilate ionospheric observations in Near-Real-Time (NRT). NRT operation naturally restricts which data sources will be available for the assimilation, as outlined in Section 2. These observations are used to produce an updated 3D representation of ionospheric electron density above 45° magnetic latitude. The assimilation runs hourly, producing outputs with a 5-min time resolution that begin 3 hr before real time. A-CHAIM also produces a simple persistence-based forecast that runs 2 hr ahead of real time. To meet the computational constraints of NRT operation, and to facilitate distribution of the output files, A-CHAIM is constructed as a series of spherical cap harmonic perturbations on E-CHAIM. This highly nonlinear state precludes the use of more traditional assimilation techniques, resulting in this first use of a particle filter for operational ionospheric modeling as described in Section 3. To assess the reliability of A-CHAIM, and of this novel application of particle filtering, the performance of the assimilation both in near-real-time and as a forecast is presented in Section 4.

2. Near-Real-Time Data

A-CHAIM must be able to take advantage of as many data sources and instrument types as possible. Any instruments that make their data available with a delay greater than a few hours will not provide much use in this context, as the ionosphere often responds to external drivers on timescales on the order of minutes, with little information retained on timescales greater than an hour (Chartier et al., 2016). The limited availability of NRT ionospheric observations is the most important consideration in the design of A-CHAIM.

2.1. Ground-Based GNSS Data

Ground-based Global Navigation Satellite System (GNSS) receivers are by far the most numerous sources of ionospheric data, with several orders of magnitude more GNSS stations providing publicly-available data than any other class of instrument. GNSS receivers are able to determine the path integrated electron density of the ionosphere between the satellite and receiver (sTEC), usually expressed in TEC Units $1 \times 10^{16} m^{-3}$ (TECu). Currently, A-CHAIM only uses data from the Global Positioning System (GPS) constellation. With each receiver able to observe 6–12 GPS satellites at a given time, a single receiver can cover a significant spatial area. With their high availability, low latency, and wide spatial coverage, GNSS TEC observations are an ideal inclusion in an NRT ionospheric data assimilation.

To extract TEC from GNSS data, one takes advantage of the dispersive propagation of radio waves in the UHF band used by GNSS, whereby ionospheric group delays and phase advances are dependent on the signal

Table 1
Data Sources Providing Ground-Based GNSS Measurements in Near-Real-Time Used by A-CHAIM

Network	Source	Source link
IGS, EUREF, GREF	German Federal Agency for Cartography and Geodesy (BKG)	igs.bkg.bund.de
CHAIN	Canadian High Arctic Ionospheric Network (CHAIN)	chain.physics.unb.ca
CDDIS	Crustal Dynamics Data Information System (CDDIS)	cddis.gsfc.nasa.gov
NOAA	NOAA National Geodetic Survey (NGS)	geodesy.noaa.gov
GARNER	California Spatial Reference Center (CSRC)	garner.ucsd.edu
CACS	Natural Resources Canada (NRCan)	rtopsdata1.geod.nrcan.gc.ca
QGN	Ministère de l'Énergie et des Ressources naturelles (MERN)	ftp.mrn.gouv.qc.ca

frequency. Using a geometry-free combination of the phase and code observables recorded on each GNSS carrier frequency, whereby the observables from each frequency are simply differenced to remove non-dispersive effects, the TEC can (to a first-order approximation) be related to the observables by

$$ sTEC = \frac{1}{A}\left(\frac{f_1^2 f_2^2}{f_1^2 - f_2^2} \right)(\Delta\phi + W + DCB_{rcv} + DCB_{sat}) \qquad (1) $$

where $A = 40.3$, f_m is the mth frequency, $\Delta\phi$ is the difference in the signal carrier phases, DCB_{rcv} and DCB_{sat} are the receiver and satellite differential code biases caused by instrumental delays, and W is a phase-leveling term used to correct an integer ambiguity in the phase-derived TEC using the code observables (Themens et al., 2013). In our use case, W is an average of the difference in the geometry-free code and phase observables over each signal lock arc, weighted by the sine of the satellite elevation (Carrano & Groves, 2009). To quality control the phase leveling process against potentially overlooked cycle slips, multipath, or leveling issues due to insufficient lock time, a standard deviation (σ) of the difference in the phase- and code-derived TEC is also recorded for each lock arc (Carrano & Groves, 2009). Any arc with $\sigma > 4.5$ TECu is discarded from the system.

In A-CHAIM, GNSS data is downloaded in the Receiver Independent Exchange Format (RINEX) format from eight different sources, listed in Table 1. The geographic distribution of these GNSS receivers is plotted in Figure 1. The downloaded files are then passed to a processing routine that converts the GNSS observables to biased TEC. The TEC data are then corrected for the satellite differential code bias DCB_{sat} using the values provided by the Institute of Geodesy and Geophysics (IGG) of the Chinese Academy of Sciences (CAS) (Wang et al., 2016). The data is not corrected for the receiver bias DCB_{rcv} in this preprocessing stage, rather the DCB_{rcv} are derived as a part of the assimilation. This is necessary in order to be able to use data from GNSS stations which do not have known DCB_{rcv}, which is the overwhelming majority of stations. Solving for the DCB_{rcv} does require additional complexity, however a full analysis of this component of the assimilation is outside the scope of this work (Table 2).

Table 2
Dependence of Modeled Observation Error R With Magnetic Latitude, as Used in A-CHAIM

Ionosonde	Observation error (R)		
Characteristic	(R_0) MLAT $\leq 45°$	MLAT $60°$	MLAT $\geq 75°$
foF2	0.15 MHz	0.30 MHz	0.45 MHz
foF1	0.25 MHz	0.50 MHz	0.75 MHz
hmF2	15 km	30 km	45 km
B0/B1→HBot	0.4 HBot	0.4 HBot	0.4 HBot

Note. Errors are smallest at the lower boundary where $R = R_0$, increasing to a maximum of $R = 3R_0$ above 75° MLAT. HBot measurements derived from B0/B1 do not vary with magnetic latitude.

2.2. Ionosondes

Ionosondes are vertically-sounding HF radars capable of providing the vertical electron density profile up to the height (hmF2) of the peak density (NmF2) of the ionosphere. These instruments have been used for ionospheric specification since the discovery of the ionosphere and formed an important component of the data set used to build E-CHAIM (Themens, Jayachandran, & Varney, 2018; Themens, Jayachandran, McCaffrey, Reid, & Varney, 2019; Themens, Jayachandran, & McCaffrey, 2019; Themens et al., 2017).

First the NOAA National Centers for Environmental Information (NCEI) is polled for available data, before the Global Ionospheric Radio Observatory (GIRO) (Reinisch & Galkin, 2011) is subsequently polled for any stations that were not available from the NOAA repository. The redundancy provided by the NOAA repository is a substantial benefit in limiting the effects of
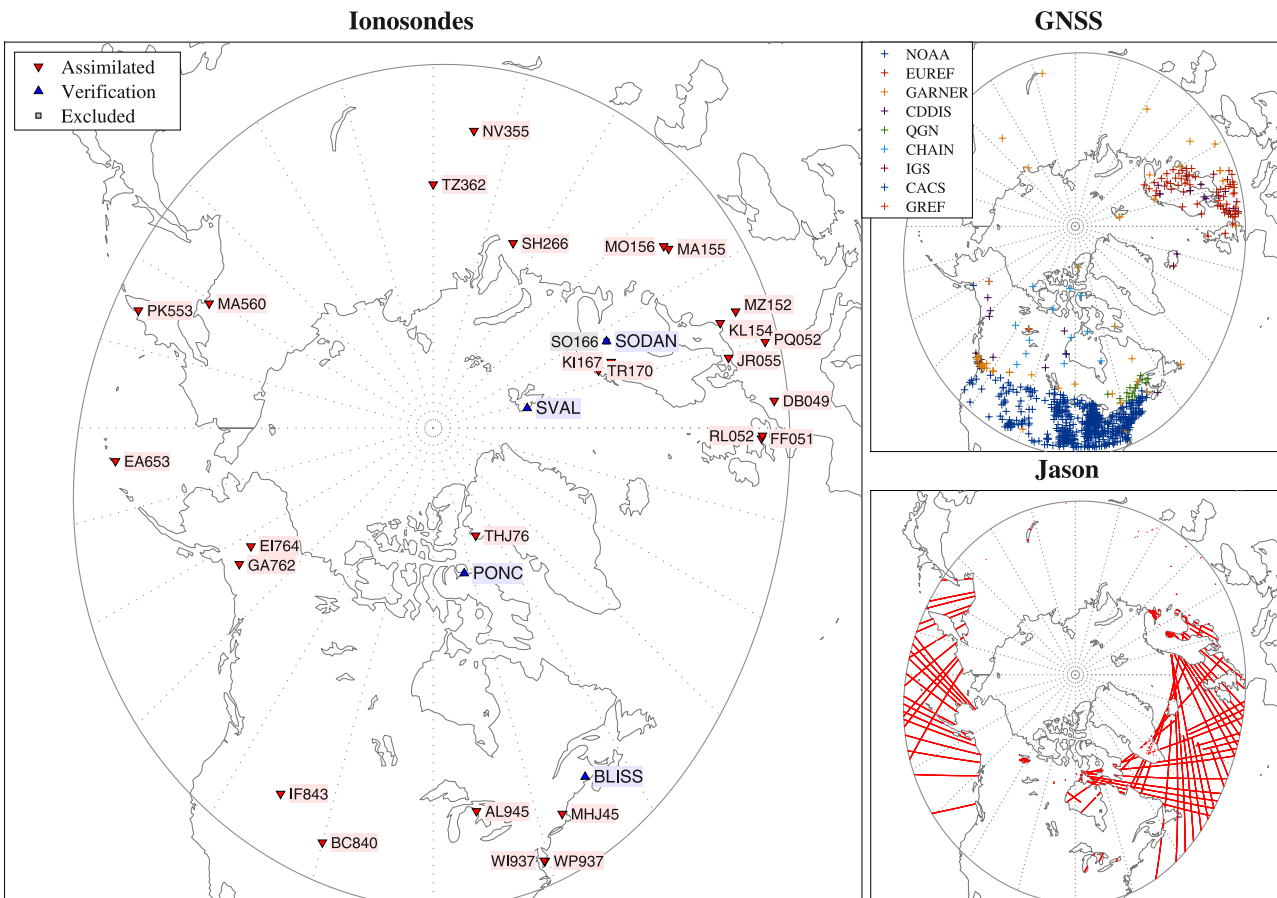
Figure 1. Geographic distribution of assimilated data sources used in A-CHAIM from September 21st though 29 September 2021. Also included are the four unassimilated ionosondes indicated in blue. The JASON figure shows all data points that were captured with a low enough latency to be included through the entire study period. The SO166 ionosonde was excluded during the test period.

service interruptions and reduces the network burden placed on any one data source. The locations of the ionosondes are noted in Figure 1. Ionograms aggregated by GIRO are processed first at the ionosonde using the local version of the ARTIST automatic ionogram scaling software (Huang & Reinisch, 2001) before both the ionogram data and processed profiles are sent to the GIRO repository. A minority of stations, such as those operated by Roshydromet, are processed using the Autoscala software with only processed profile information being sent to the GIRO database.

A-CHAIM uses five of the ionosonde-derived characteristics: $foF2$, $hmF2$, $foF1$, $B0$, and $B1$. Each characteristic is treated as being independent, even when they are derived from the same ionogram. $B0$ and $B1$ are parameters that control the bottomside profile in the IRI (Altadill et al., 2009), and are converted to the E-CHAIM equivalent H_{Bot} (Section 3.1) using a nonlinear fit to the equivalent IRI shape. Autoscaled ionosonde measurements do not have Gaussian measurement errors, and the true error varies widely with instrument latitude, geomagnetic and solar activity, interference environment, instrument design and configuration, as well as the version and configuration of the autoscaling software. Lacking a universal method to determine the errors, A-CHAIM uses the following simple heuristic to generate Gaussian errors. All characteristics except H_{Bot} have observation errors that increase with magnetic latitude, to reflect both the greater likelihood of scaling errors, and the variability within the assimilation window. This error R is modeled as a minimum error R_0 scaled by a simple transition function (2).

$$R = R_0(2 + \tanh((mlat - 60^o)/5)) \tag{2}$$

At the lower boundary $R = R_0$, increasing with magnetic latitude through the auroral region to peak at $R = 3R_0$ in the polar cap. In addition, a filter on $hmF2$ is applied to the data, where all characteristics from a sounding with

*hmF*2 < 175 km or *hmF*2 > 450 km are rejected. This simple filter catches highly biased ionogram scaling errors typically associated with scaler early stopping; where the scaling routine truncates a trace prematurely, or with the scaler missing the F2 trace and misinterpreting the F1 layer as the F2 layer (Themens et al., 2022).

2.3. Altimeter

A-CHAIM also makes use of space-borne altimeter data from the JASON-3 satellite mission, provided by the NOAA National Oceanographic Data Center. As a by-product of the altimeter solution for sea-surface height, vertical ionospheric TEC above the ocean can be inferred (Li et al., 2018). This is done following the same concept as GNSS TEC products, where JASON's Ku band antenna excess phase can be directly related to the TEC along the ray path as

$$vTEC = \frac{dRf^2}{40.3} \tag{3}$$

where dR is the excess ground range and f = 13.575 GHz is the signal frequency. The resulting TEC is then filtered to remove outliers and ground/ice scatter using the provided quality flags. While the overall precision of the JASON TEC is 4 TECU, it is largely unbiased and provides a crucial constraint over the oceans, where no other data set has adequate coverage.

2.4. Latency

The A-CHAIM system has been gathering data and running in real time since 2020. During this time, it has been updated several times. To assess the performance of the system, we will focus on the period from September 21st through 29 September 2021. This time period includes a moderate Kp 4 event, with a M2.8 solar flare on September 23rd, providing an opportunity to study how the assimilation behaves during disturbed conditions. The results in this study were generated with the latest version of A-CHAIM in an offline run. To ensure that there were minimal difference between the offline run and the real-time performance, the results presented here were generated using the actual data collected during each hour by the online system, as well as the outputs of the background model E-CHAIM as they were produced in real time.

E-CHAIM uses several geophysical indices to produce a storm model (Themens et al., 2017). However, when operating in near-real-time, or producing a forecast, these indices may not be available, and therefore the storm model cannot operate as normal. The performance of the background model E-CHAIM is therefore dependent on the time when the model was run. In general, the storm model will turn off at an unpredictable time mid-run, which is dependent on the specific timings of the index providers. This race condition behavior is hard to model outside of a real-time setting, and is the reason why we use the outputs of E-CHAIM that were generated in real time.

At each hour, A-CHAIM takes the data file produced from the processing pipeline and begins a run. This run nominally starts 3 hours in the past, using the output of the previous hour's run to provide the initial conditions. A-CHAIM proceeds forward through the present time, and continues until it reaches 2 hr into the future. This provides a low-skill persistence forecast. Figure 2 shows the number of observations from each instrument type available during each 5 min assimilation window. It is clear that the amount of available data varies greatly depending on how close the assimilation time is to the present time.

Through successive runs of A-CHAIM, the same assimilation window is traversed 5 times, twice as a forecast, and thrice with actual data. Each output A-CHAIM generates for that assimilation window was produced with very different amounts of data, and so it is important to assess the performance of each of these five versions of the same time window. We can group these separated by the number of hours latency from the time when the data was collated and passed to the assimilation routine. These are labeled t−02h, t−01h, t−00h, t+00h, and t+01h, indicating the number of hours from the current time, rounding toward zero. Each of these sets of a given latency forms a continuous time series of outputs, and each latency contains every assimilation window exactly once. We can therefore compare the performance of each of these latencies to each other.

Throughout the study period, the number of instruments of each type were recorded for each 5-min assimilation window. The results are summarized in Figure 3 for each instrument type, separated by latency. The form of these
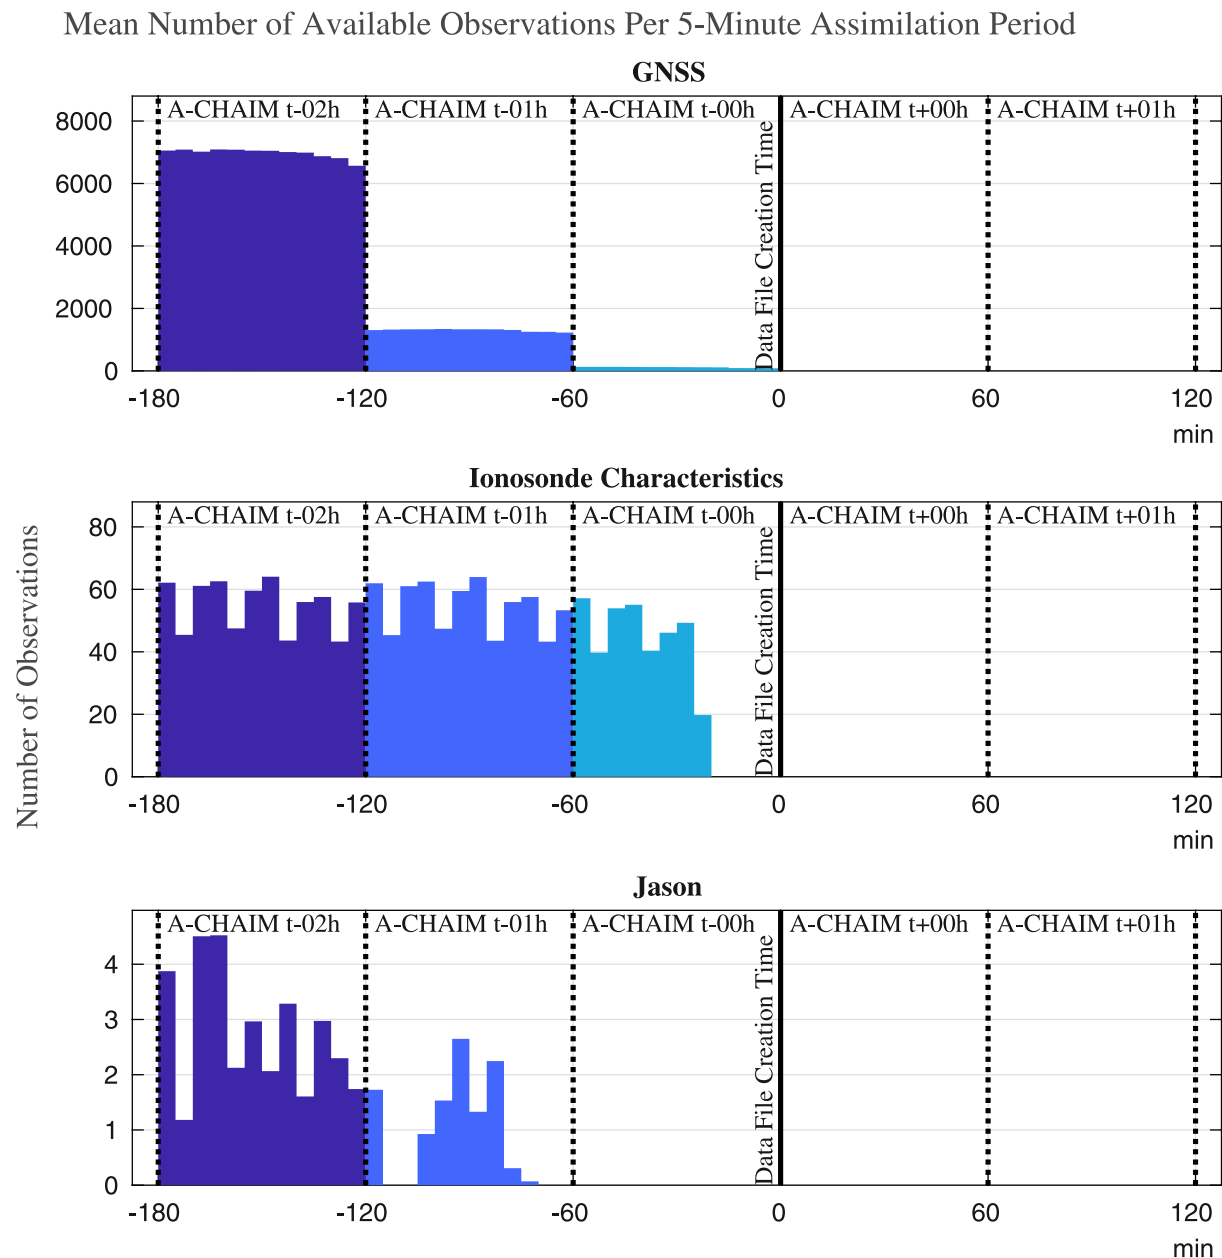
Mean Number of Available Observations Per 5-Minute Assimilation Period

Figure 2. Histogram showing the mean number of observations available per instrument type, relative to the current time of the assimilation. Each hour a data file is created containing all observations from the past 3 hr. A new run of A-CHAIM is initialized from the previous hour's run, and then assimilates all of the data in this file. As latency has a large effect on the number of observations available, A-CHAIM outputs are labeled based on the number of hours removed from the time of data file.

plots is markedly different for each data source. For GNSS receivers, as we pass the threshold of each hour, we lose a significant fraction of the number of stations reporting data. At t−02h, we have an average of 537 stations reporting, at t−01h this drops to 121, and at t−00h only 15 stations are reporting. As most NRT GNSS data is distributed in the form of hourly RINEX files, no data is available until the hour has finished. As a result, the number of GNSS observations does not change significantly during any given hour. Accordingly, there would be little benefit to running A-CHAIM more frequently than hourly. While very low latency GNSS data is also available using Network Transport of RTCM via Internet Protocol (NTRIP), few networks fully implement NTRIP at the present time and particularly few with coverage at remote high latitudes. At the present time NTRIP is not used by A-CHAIM, but is planned for future implementation.
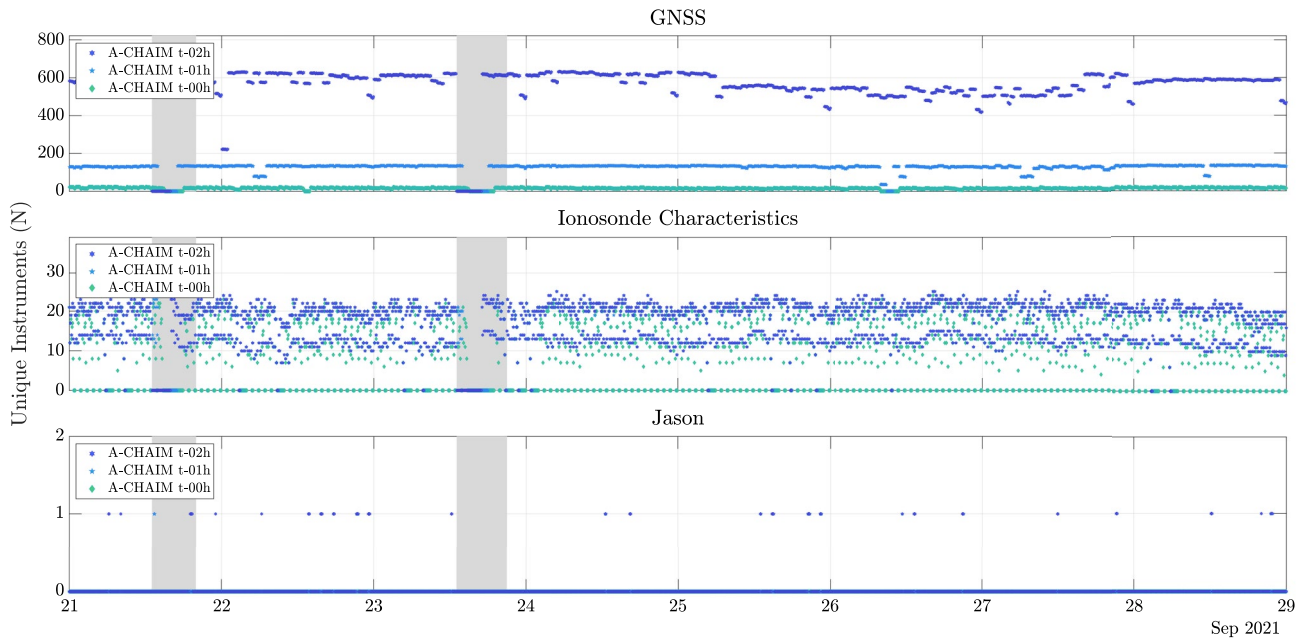
Figure 3. Number of unique instruments reporting data at each latency during the study period, 21 September through 29 September 2021.

Ionosondes, which are distributed as individual files for each sounding, do not show this sharp hourly transition. The number of ionosondes reporting data is nearly identical for both t−02h and t−01h. There is less data available at only t−00h, but the average number of stations available is still half that at t−02h. After a single hour, nearly all ionosonde data that will be available has already been published. There is a distinctive double-sinusoidal pattern to the ionosonde latency data, which is an aliasing effect of the sounding schedules of the individual ionosondes.

The intermittent nature of JASON data is clearly visible in Figure 2. JASON data is distributed in files that cover discrete lengths of time, and so like GNSS data the data from the beginning of a file period is only available once the file has ended. Accordingly, JASON data is only useable when the high-latitude measurements happen to fall near the end of the observation file. When data is available, JASON provides hundreds of measurements over a broad area, and so provides a useful NRT data source despite its intermittent nature.

3. Assimilation Method

The objective of data assimilation is generally to produce a statistical measure of the behavior of a system, given a certain set of observations and model constraints. The relatively small state space used in A-CHAIM comes at the cost of unavoidable nonlinearity, which necessitates the use of a particle filtering technique. As particle filters are still relatively unknown in the ionospheric physics community, a basic outline is given in Section 3.2. A more complete treatment of particle filters is given in the very accessible Doucet and Johansen (2009). Particle filters are part of a broad class of statistical models known as Hidden State Markov models. This class includes many other techniques, like the ubiquitous Kalman filters that are often used in data assimilation (Doucet et al., 2000).

3.1. Parameterization of the Ionospheric Electron Density

E-CHAIM is parameterized as a set of ionospheric profile parameters, used to reconstruct the vertical electron density profile via a semi-Epstein layer formulation. These parameters are specified as coefficients which depend on the date, time, and geophysical indices, expanded horizontally in terms of spherical cap harmonics in Altitude Adjusted Corrected Geomagnetic Coordinates (AACGM) coordinates (Shepherd, 2014). A full description of the relevant components of the E-CHAIM parameterization can also be found in Themens et al. (2017), Themens, Jayachandran, and Varney (2018), and Themens, Jayachandran, and McCaffrey (2019). In A-CHAIM we use the same vertical parameterization as the E-CHAIM model. The shape of the profile is controlled by eight parameters, $NmF2$, $hmF2$, $hmF1$, hmE, H_{Bot}, H_{Top}, H_{F1}, and H_E. Electron density for a height h is given by:

$$N_e(h) := NmF2 \cdot \text{sech}^2\left(\frac{h - hmF2}{H(h)}\right) \tag{4}$$

$$H(h) := \begin{cases} 2H_{Top} \cdot \left(1 + \dfrac{rg(h - hmF2)}{rH_{Top} + g(h - hmF2)}\right) & h \geq hmF2 \\[2em] H_B(h) \cdot \left(\dfrac{1}{1 + \exp\left(\dfrac{HmE - 15 - h}{2.5}\right)}\right) & h < hmF2 \end{cases} \tag{5}$$

$$H_B = H_{Bot} + H_{F1} \cdot \text{sech}^2\left(\frac{h - hmF1}{(hmF2 - hmF1)/2.5}\right) + H_E \cdot \text{sech}^2\left(\frac{h - hmE}{25}\right) \tag{6}$$

$g = 0.18$ and $r = 20$ are constants (Themens, Jayachandran, Bilitza, et al., 2018).

In E-CHAIM, auroral electron precipitation is represented by a semi-physical precipitation scheme outlined in Watson et al. (2021). While this module performs well, it can be computationally intensive to calculate for ensemble assimilation methods; as such, in A-CHAIM an additional layer is used to represent the electron density enhancement from precipitating electrons. This is modeled as a Chapman function with height-varying scale height, which is described by an additional four parameters: NmP, hmP, H_1P, and H_2P.

$$N_e(h) = NmP \cdot \exp(1 - z - \exp(-z)) \tag{7}$$

$$z = \frac{h - hmP}{H_1P - H_2P \cdot \frac{\exp(-(h-hmP)/15)}{1+\exp(-(h-hmP)/15)}} \tag{8}$$

The parameters for the auroral layer are produced by fitting the Chapman layer to a 3D gridded output of the E-CHAIM precipitation module. Without this layer, the vertical parameterization would not be able to describe auroral electron densities, which can be a significant contributor to observed TEC. This would result in decreased performance during strong auroral enhancements, as the increase in TEC could be falsely attributed to some other feature, such as an F-region enhancement.

The 12 parameters above provide a specification of the vertical electron density profile above a single geographic point, at a single time. To describe the geographic variation in these vertical profile parameters, as shown in Figure 4, the output of the E-CHAIM model is fitted with a spherical cap harmonic expansion in centered dipole (CD) coordinates. This basis is defined in a cone with half-angle $\theta_c = 51°$ centered on the North geomagnetic pole. The CD colatitude is rescaled as in Equation 9 so that the south pole of the regular spherical harmonic basis is coincident with the cap. For any profile parameter f, with rescaled magnetic dipole colatitude θ and CD geomagnetic longitude ϕ:

$$\theta = \frac{180^o}{\theta_c}(90^o - mlat_{CD}) \tag{9}$$

$$f(\theta, \phi) = \sum_{l=0}^{\infty}\sum_{m=0}^{l} Y_{lm} = \sum_{l=0}^{\infty}\sum_{m=0}^{l} P_{lm}(\cos\theta)(C_{lm}\cos m\phi + S_{lm}\sin m\phi) \tag{10}$$

The coefficients of this expansion, C_{lm} and S_{lm}, specify the 2D geographic variation of the corresponding profile parameter f. Together all of these maps of profile parameters specify a complete 3D ionosphere state at a fixed time. This parameterization has several notable advantages. The shape of the electron density profile is constrained to be physically realistic, unlike parameterizations that use discrete points at fixed altitudes. The electron density is also guaranteed to be spatially smooth and differentiable for raytracing applications. It also allows a complete description of the 3D electron density with relatively few parameters. Using 12 orders of spherical cap harmonics for each vertical profile parameter, the entire state can be specified with only 1,352 parameters, with an additional 676 for the auroral precipitation.

The size of the state can be further reduced by removing certain parameters from the assimilation. The contributions to electron density from the E and F1 layers outside of the auroral region are both relatively well captured
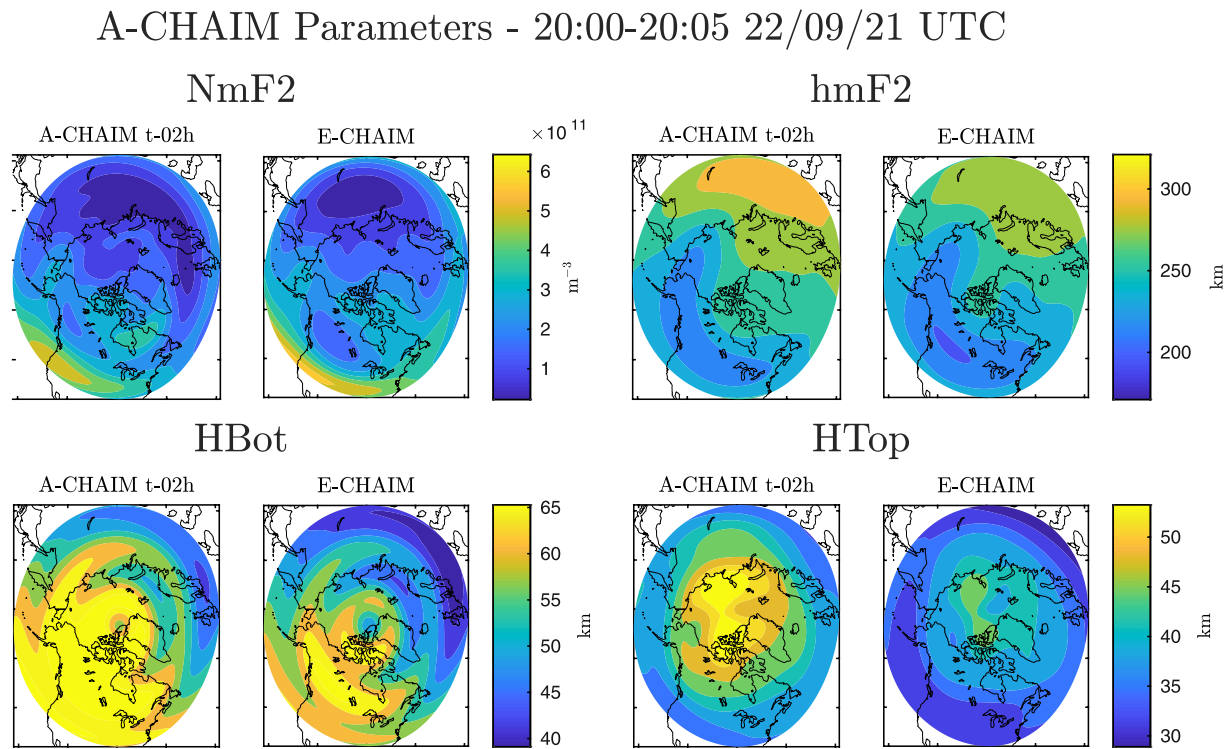
A-CHAIM Parameters - 20:00-20:05 22/09/21 UTC

Figure 4. Example of the geographic variation of the four active ionospheric profile parameters at 22:00 UTC, 22 September 2021. All other parameters are not updated by the assimilation, and so are identical for both models.

by empirical models, being driven primarily by solar activity. Additionally, variations in these layers do not contribute significantly to TEC and so most of the available data is not sensitive to these parameters. Rather than updating the E, F1, or auroral layers in the assimilation, we can keep the corresponding parameters fixed to their empirical values. By excluding $hmF1$, hmE, H_{F1}, H_E, NmP, hmP, H_1P, and H_2P, and preserving $NmF2$, $hmF2$, H_{Bot}, H_{Top} we can further reduce the number of parameters to estimate to only 676.

Each set of coefficients describes the configuration of the ionosphere for a single fixed time. To capture the time evolution of the system, it is necessary to specify a set of these coefficients for each 5-min assimilation window. The amount of data available during any one assimilation window is usually insufficient to produce a well-determined solution for the optimal set of coefficients, and therefore unlikely to produce a good estimate of the ionospheric state. To be able to produce a physically realistic model, information from one time period must be passed forward to subsequent times. The most probable history of model coefficients, and therefore ionospheric states, must be determined.

3.2. Particle Filters

Any given parametrization of our system defines vector space $\mathbb{X}$, the set of all possible configurations of system. A specific configuration of the system is described by a state vector $\mathbf{x} \in \mathbb{X}$. For the purposes of this model we will treat $\mathbf{x}$ as describing an instantaneous and static configuration of the system. The system is assumed to maintain this fixed state $\mathbf{x}_k$ for a discrete period of time δt at a time t_k. During each time interval t_k we record some set of observations $\mathbf{y}_k$. The time evolution of the system can then be modeled as a succession of states $\mathbf{x}_{1:n}$ at discrete times $t_{1:n}$. Naturally, the true configuration of a system like the ionosphere is not directly observable. The trajectory in state space $\mathbf{x}_{1:n}$ is hidden, and only indirectly measured through the observations $\mathbf{y}_{1:n}$.

$$t_{1:n} = \{t_1, t_2, \ldots, t_{n-1}, t_n\}$$
$$\mathbf{x}_{1:n} = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{n-1}, \mathbf{x}_n\} \qquad (11)$$
$$\mathbf{y}_{1:n} = \{\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_{n-1}, \mathbf{y}_n\}$$

The observations $\mathbf{y}_k$ are subject to error and so are a random variable sampled with a likelihood $p(\mathbf{y}_k|\mathbf{x}_k)$. To restate the problem of data assimilation in these terms, we seek to make some inference about the hidden trajectory of our system in state space $\mathbf{x}_{1:n}$ given some set of imperfect observations $\mathbf{y}_{1:n}$. Using Bayes' theorem, this requires evaluating the following expressions.

$$p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n}) = \frac{p(\mathbf{x}_{1:n})p(\mathbf{y}_{1:n}|\mathbf{x}_{1:n})}{p(\mathbf{y}_{1:n})}$$

$$p(\mathbf{y}_{1:n}) = \int p(\mathbf{x}_{1:n})p(\mathbf{y}_{1:n}|\mathbf{x}_{1:n})d\mathbf{x}_{1:n}$$

(12)

The probability of moving from a state $\mathbf{x}_{n-1}$ to $\mathbf{x}_n$ is given by the transition probability, or forecast model $f(\mathbf{x}_n|\mathbf{x}_{n-1})$.

As we do not have a physics-based model to perform the forecasting step of the assimilation, we must use some other method to propagate the state forward in time. If we let $\mathbf{u}_n \in \mathbb{X}$ be the state vector corresponding to the background model at time t_n, then we can simply propagate the state forward by following the movement in the background model from $\mathbf{u}_{n-1}$ to $\mathbf{u}_n$. While this is clearly unable to capture the dynamic evolution of the ionosphere, it will at least allow the assimilation to follow the diurnal variations that are well described by climatology. We would also like the state to gradually converge with the background model over time, which can be controlled with a multiplicative factor $\lambda \in [0, 1]$. In A-CHAIM $\lambda = 0.95$ is used for all parameters, which causes the distance between the particles and the background to approximately halve every hour. The non-stochastic components of the time evolution of our model define a forecast operator $\mathbf{F}_n(\mathbf{x}_{n-1})$ Equation 13.

$$\mathbf{F}_n(\mathbf{x}_{n-1}) = \lambda(\mathbf{x}_{n-1} + \mathbf{u}_n - \mathbf{u}_{n-1}) + (1 - \lambda)\mathbf{u}_n$$

(13)

To compensate for our incomplete knowledge of the processes involved, there is some uncertainty in the time evolution of the ionospheric system. This stochastic behavior can be modeled by assuming that the ionospheric state will deviate from Equation 13 by some random deviation drawn from a multivariate Gaussian distribution $\mathcal{N}(0, \mathbf{Q}_n)$. $\mathbf{Q}_n$ is a positive definite matrix, analogous to the noise covariance matrix in a Kalman filter. The precise form of $\mathbf{Q}_n$ is discussed in detail in 3.5 Combining this noise process with our forecast operator Equation 13 gives us the transition probability, $f(\mathbf{x}_n|\mathbf{x}_{n-1})$, the probability of moving from a state $\mathbf{x}_{n-1}$ to $\mathbf{x}_n$.

$$f(\mathbf{x}_n|\mathbf{x}_{n-1}) = \exp\left[-\frac{1}{2}(\mathbf{x}_n - \mathbf{F}_n(\mathbf{x}_{n-1})^T\mathbf{Q}_n^{-1}(\mathbf{x}_n - \mathbf{F}_n(\mathbf{x}_{n-1})\right]$$

(14)

This forecast model also defines the a priori probability density of the state space $p(\mathbf{x}_{1:n})$:

$$p(\mathbf{x}_{1:n}) = \prod_{k=1}^{n} f(\mathbf{x}_k|\mathbf{x}_{k-1})$$

(15)

Particle filters are a Monte Carlo technique that use an ensemble of samples $\mathbf{X}_{1:n} \in \mathbb{X}$ as an empirical approximation $\hat{p}(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$ of a target density $p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$. Using the Dirac delta distribution δ:

$$\mathbf{X}_{1:n} \sim p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$$

(16)

$$\hat{p}(\mathbf{x}_{1:n}|\mathbf{y}_{1:n}) = \sum_{i=1}^{N} \delta_{\mathbf{X}_{1:n}^i}(\mathbf{x}_{1:n})$$

(17)

As $p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$ is the unknown density which we wish to approximate, it is not practical to sample from it directly. Instead, we sample from a more tractable importance density $q(\mathbf{x}_{1:n})$ that has the same support. The exact form of $q(\mathbf{x}_{1:n})$ can vary depending on the objective. Generally the best performance is maintained by keeping the importance density as close to $q(\mathbf{x}_n|\mathbf{y}_{1:n}, \mathbf{x}_{1:n-1}) = p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n}, \mathbf{x}_{1:n-1})$ as is practical (Doucet & Johansen, 2009). We are able produce an ensemble $\mathbf{X}_{1:n} \sim q(\mathbf{x}_{1:n})$, hereafter called particles, which allow us to reconstruct the original density $p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$ by assigning each particle a weight. Each particle $\mathbf{X}_{1:n}^i$ has an unnormalized weight $w_n(\mathbf{X}_{1:n}^i)$ at time t_n given by:

$$w_1(\mathbf{x}_1) = \frac{p(\mathbf{x}_1)p(\mathbf{y}_1|\mathbf{x}_1)}{q(\mathbf{x}_1)}, \quad w_n(\mathbf{x}_{1:n}) = w_1(\mathbf{x}_1)\prod_{k=2}^{n}\frac{f(\mathbf{x}_k|\mathbf{x}_{k-1})p(\mathbf{y}_k|\mathbf{x}_k)}{q_k(\mathbf{x}_k|\mathbf{x}_{k-1})}$$

(18)

when we normalize the weights $w_n(\mathbf{x}_{1:n})$ our sum in Equation 17 takes the following form:

$$\mathbf{X}_{1:n} \sim q(\mathbf{x}_{1:n}|\mathbf{y}_{1:n}) \tag{19}$$

$$W_n^i = \frac{w_n(\mathbf{X}_{1:n}^i)}{\sum_{j=1}^N w_n(\mathbf{X}_{1:n}^j)} \tag{20}$$

$$p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n}) \approx \hat{p}(\mathbf{x}_{1:n}|\mathbf{y}_{1:n}) = \sum_{i=1}^N W_n^i \delta_{\mathbf{X}_{1:n}^i}(\mathbf{x}_{1:n}) \tag{21}$$

By definition, the particles $\mathbf{X}_{1:n}^i$ are sampled from the importance density $q(\mathbf{x})$. Using the forecast model $f(\mathbf{x}_n|\mathbf{x}_{n-1})$, and the observations $\mathbf{y}_{1:n}$, we are able to generate weights w_n^i for each particle. These weights, once normalized, allow us to estimate the probability density we are interested in, $p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$ using Equation 21. We can also take the expectation value of a function $\phi(\mathbf{x}_{1:n})$, allowing the calculation of any statistical moment of interest.

$$\langle \phi(\mathbf{x}_{1:n}) \rangle = \int \phi(\mathbf{x}_{1:n}) p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n}) d\mathbf{x}_{1:n} \approx \frac{1}{N} \sum_{i=1}^N W_n^i \phi\left(\delta_{\mathbf{X}_{1:n}^i}\right) \tag{22}$$

Critically for A-CHAIM, Equation 22 allows the calculation of the expectation value of electron density, $\langle N_e(\mathbf{x}_{1:n}) \rangle$, using Equation 4.

3.3. Weight Degeneracy and Resampling

The weights w_n^i provide a measure of how probable a particle $\mathbf{X}_{1:n}^i$ is, given the observations $\mathbf{y}_{1:n}$. Higher weight particles are more likely, and so contribute more to the weighted sum in Equation 22. Conversely, lower weight particles contribute less. As n increases, the variance of the estimates produced by the set of particles $\mathbf{X}_{1:n}$ tends to increase dramatically. The unnormalized weight of a particle $w_n(\mathbf{X}_{1:n}^i) \propto \prod_{k=1}^n p(\mathbf{y}_k|\mathbf{X}_k^i)$, so even small differences between particles become magnified over time. The near-inevitable result is weight degeneracy, where only a single particle will have a non-zero weight, and all other particles having an identical and permanent weight of zero (Bengtsson et al., 2008). In order to prevent this issue, sequential sampling schemes need to re-generate their particles through a process known as resampling.

In resampling we use our weighted ensemble of particles $\mathbf{X}_{1:n}^i \sim q(\mathbf{x}_{1:n})$ to produce an unweighted ensemble of particles $\tilde{\mathbf{X}}_{1:n}^i \sim p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$. This is simple to accomplish by taking a random sample from $\mathbf{X}_{1:n}^i$ with probability $W_{1:n}^i$. There are several unbiased resampling methods in the particle filter literature, and we have used the simple and common method known as systematic resampling (Douc et al., 2005). This produces another Monte Carlo approximation to $p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$

$$\tilde{\mathbf{X}}_{1:n}^i \sim p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n}) \tag{23}$$

$$p(\mathbf{x}_{1:n}|\mathbf{y}_{1:n}) \approx \tilde{p}(\mathbf{x}_{1:n}|\mathbf{y}_{1:n}) = \sum_{i=1}^N \frac{1}{N} \delta_{\tilde{\mathbf{X}}_{1:n}^i}(\mathbf{x}_{1:n}) \tag{24}$$

This is equivalent to a weighted sum where all of the weights are $\frac{1}{N}$. By replacing our original sample $\mathbf{X}_{1:n}^i$ with $\tilde{\mathbf{X}}_{1:n}^i$, and $W_{1:n}^i$ with $\frac{1}{N}$ we have reset our particle weights and prevented weight degeneracy. Resampling will tend to remove low-weight particles from $\mathbf{X}_{1:n}^i$ and replace them with copies of high-weight particles, at the cost of some loss of information. It is therefore optimal to use the weighted particles to calculate statistical moments of interest from $\hat{p}(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$ before resampling.

At a fundamental level, particle degeneracy is a result of the ever-increasing dimensionality of the particle trajectories through state space. At each time t_n, the number of dimensions occupied by $\mathbf{X}_{1:n}^i$ increases by the size of $\mathbb{X}$ (Doucet & Johansen, 2009), while the number of particles remains fixed. While resampling is able to alleviate the issues created by increasing dimensionality over time, this becomes a more critical problem in inherently high-dimensional particle filters. For large scale geophysical systems, the size of the state space is great enough that no realistic number of particles can prevent degeneracy in a basic particle filter (Bengtsson et al., 2008). van Leeuwen et al. (2019) provides an in-depth review of more sophisticated particle filtering tech-

niques implemented across geophysics to avoid this dimensionality issue. Localization is given particular focus, as it allows for the separate treatment of small subsets of the state and data spaces. This dramatically reduces the dimensionality of the problem. Other data assimilation techniques also use localization, such as the Local Ensemble Transform Kalman Filter (LETKF) (Ott et al., 2004). One major drawback of localization for ionospheric studies is that GNSS sTEC, the most numerous and widespread electron density measurement, is inherently non-local. As the electron density is integrated along line-of-sight, there is no specific point in space where the observation took place. Using non-local observations in a LETKF is still an open problem in data assimilation (van Leeuwen, 2019). GNSS sTEC measurements are non-localizable, non-linear, highly correlated, biased, and have very low information content on a per-measurement basis. Their utility in assimilation is only due to the very large number of observations available. In order to perform adequately given the operational constraints of A-CHAIM, a new solution to particle degeneracy had to be developed.

3.4. Composite Observations

Without access to localization, some other approach must be taken to reduce the dimensionality of the problem. To accomplish this, we will combine the real observations into a new, composite observable process, which exists in a lower-dimensional observable space. With a few caveats, this composite observable space preserves all of the desired properties of the original observations.

For clarity of notation we will consider only a single time t_n. Our measurements $\mathbf{y}_n$ form an m–dimensional vector of observables by which we infer the hidden process $\mathbf{x}_n$. If there exists some subset of our measurements that are independent from any measurement outside of the subset, then it is possible to factor the likelihood $p(\mathbf{y}_n|\mathbf{x}_n)$. Let $\iota \subset \{1, 2, \ldots, m\}$ be an index set so that $\mathbf{y}_n^\iota \perp \mathbf{y}_n^{\neg\iota}$, then $p(\mathbf{y}_n|\mathbf{x}_n) = p(\mathbf{y}_n^\iota|\mathbf{x}_n)p(\mathbf{y}_n^{\neg\iota}|\mathbf{x}_n)$. Let $l_n^\iota(\mathbf{x}_n, \mathbf{y}_n^\iota)$ be a function of the log likelihood of the observations indexed by ι.

$$l_n^\iota = l_n^\iota(\mathbf{x}_n, \mathbf{y}_n^\iota) = -2\log p(\mathbf{y}_n^\iota|\mathbf{x}_n) \tag{25}$$

If $p(\mathbf{y}_n^\iota|\mathbf{x}_n)$ can be modeled as a multivariate Gaussian with an observation error covariance $\mathbf{R}_n^\iota$, then $p(l_n^\iota|\mathbf{x}_n)$ has a comparatively simple closed-form solution. Let $n(\iota)$ be the number of observations included in the subset. As the sum of $n(\iota)$ squared random variables, the likelihood of the combined observations $p(l_n^\iota|\mathbf{x}_n)$ are χ^2 distributed with $n(\iota)$ degrees of freedom. See Appendix A for a brief derivation.

$$p(l_n^\iota|\mathbf{x}_n) = \frac{1}{2^{n(\iota)/2}\Gamma(n(\iota)/2)}(l_n^\iota)^{n(\iota)/2-1}\exp(-l_n^\iota/2) \tag{26}$$

As l_n^ι is a function of $\mathbf{x}_n$, it is not an observable. It is a statistic that measures how well the residuals of $\mathbf{y}_n^\iota - \mathbf{H}^\iota(\mathbf{X}_n^\iota)$ fit the expected distribution specified by the error covariance matrix $\mathbf{R}_n^\iota$. Rather than the $n(\iota)$-dimensional $\mathbf{y}_n^\iota$, l_n^ι is a scalar.

Let $\tau_n = \left\{\iota_1, \iota_2, \ldots, \iota_\mu \mid \mathbf{y}_n = \bigcup_{\iota \in \tau_n} \mathbf{y}_n^\iota\right\}$ be a set of index sets that form a partition of $\mathbf{y}_n$. Any choice of partition τ_n must have $1 \le \mu \le m$ elements. If we choose each component of the partition $\tau_n = \left\{\iota_1, \iota_2, \ldots, \iota_\mu \mid \mathbf{y}_n^\iota \perp \mathbf{y}_n^\xi \forall \iota \ne \xi, \in \tau_n\right\}$ so that each $\mathbf{y}_n^\iota$ is independent of every other.

$$\mathbf{l}_n = \{l_n^\iota|\iota \in \tau_n\} \tag{27}$$

Using Equation 26 we can find $p(\mathbf{l}_n|\mathbf{x}_n)$.

$$p(\mathbf{l}_n|\mathbf{x}_n, \tau_n) = \prod_{\iota \in \tau} p(l_n^\iota|\mathbf{x}_n) \tag{28}$$

If we replace the target distribution in Equation 12 with Equation 29

$$p(\mathbf{x}_{1:n}|\mathbf{l}_{1:n}, \tau_{1:n}) = \frac{p(\mathbf{x}_{1:n})p(\mathbf{l}_{1:n}|\mathbf{x}_{1:n}, \tau_{1:n})}{p(\mathbf{l}_{1:n}, \tau_{1:n})}$$

$$p(\mathbf{l}_{1:n}, \tau_{1:n}) = \int p(\mathbf{x}_{1:n})p(\mathbf{l}_{1:n}|\mathbf{x}_{1:n}, \tau_{1:n})d\mathbf{x}_{1:n} \tag{29}$$

We can re-derive all of the particle filter equations in terms of this new target distribution. Rather than working in the m-dimensional space $\mathbf{Y}$, we are in the μ-dimensional space $\mathcal{Y}$. The densities in state space $q(\mathbf{x})$ and $f(\mathbf{x})$ are unchanged.

$$w_1(\mathbf{x}_1) = \frac{p(\mathbf{x}_1)p(\mathbf{l}_1|\mathbf{x}_1, \tau_1)}{q(\mathbf{x}_1)}, \quad w_n(\mathbf{x}_{1:n}) = w_1(\mathbf{x}_1)\prod_{k=2}^{n}\frac{f(\mathbf{x}_k|\mathbf{x}_{k-1})p(\mathbf{l}_k|\mathbf{x}_k, \tau_k)}{q_k(\mathbf{x}_k|\mathbf{x}_{k-1})} \tag{30}$$

$$\mathbf{X}_{1:n} \sim q(\mathbf{x}_{1:n}|\mathbf{l}_{1:n}, \tau_{1:n}) \tag{31}$$

$$W_n^i = \frac{w_n(\mathbf{X}_{1:n}^i)}{\sum_{j=1}^{N} w_n(\mathbf{X}_{1:n}^j)} \tag{32}$$

$$p(\mathbf{x}_{1:n}|\mathbf{l}_{1:n}, \tau_{1:n}) \approx \hat{p}(\mathbf{x}_{1:n}|\mathbf{l}_{1:n}, \tau_{1:n}) = \sum_{i=1}^{N} W_n^i \delta_{\mathbf{X}_{1:n}^i}(\mathbf{x}_{1:n}) \tag{33}$$

The partition of the observation space τ is not prescribed, other than being constrained by the statistical independence of the measurements. This flexibility allows tuning of the particle filter for any number of experimental objectives. For example, τ could be chosen to minimize the variation of the weights w^i. In A-CHAIM the observations are partitioned such that each instrument type is handled separately. This largely eliminates the problem of particle degeneracy, which was the primary objective. It also solves the problem of trying to balance the relative influence of each instrument type on the assimilation. The number of GNSS, ionosonde, and altimeter observations during any given assimilation step usually fluctuate by orders of magnitude, and so being able to assimilate each instrument type independently removes the need for ad-hoc solutions like kriging, spoofing or duplicating data. Partitioning and recombining the observations $\mathbf{y}_n$ into composite observables $\mathbf{l}_n$ allows us to recover many of the advantages of localization, even in a system that does not easily permit localization.

3.5. Forecast Optimization

The forecast model in A-CHAIM has a deterministic component (13) and a stochastic component. Let $\mathbf{X}_{n|n-1}^i$ be the particles at time t_n after the deterministic component of the forecast step, but before any stochastic movement is added.

$$\mathbf{X}_{n|n-1}^i = \mathbf{F}(\mathbf{X}_{n-1}^i) \tag{34}$$

The stochastic component is modeled as a Gaussian random noise as in Equation 14. At each time we add a random displacement $\delta\mathbf{X}_n^i \sim \mathcal{N}(0, \mathbf{Q}_n)$.

$$\mathbf{X}_n^i = \mathbf{X}_{n|n-1}^i + \delta\mathbf{X}_n^i \tag{35}$$

It was determined though experimentation that a minimum stochastic variance at any timestep $\mathbf{Q}_n^{\min} = \text{diag}\left(\left(\frac{\mathbf{u}_n - \mathbf{u}_{n-1}}{2}\right)^2\right)$ allowed the filter to perform well during calm conditions, but was not able to adapt quickly enough during storm periods. Choosing a fixed covariance that was able to capture storm behavior would degrade the filter performance during quiet periods. It is therefore necessary to evolve the diagonal variance matrix $\mathbf{Q}_n$ with the particle filter, to be able to adapt to changing ionospheric variability. A simple way to accomplish this without reliance on external drivers is to monitor the stochastic movements of previous timesteps. By examining the step sizes of higher weight particles, we can estimate an improved variance $\tilde{\mathbf{Q}}_n = \text{diag}(E[(\delta\mathbf{X}_n)^2])$. This variance estimate tends to be very noisy, and can tend to produce unstable behavior if not tempered. In A-CHAIM $\mathbf{Q}_n$ is updated with a simple algorithm, but more sophisticated techniques to estimate this variance are certainly possible.

$$\mathbf{Q}_0 = \mathbf{Q}_0^{\min}, \quad \mathbf{Q}_n = \lambda\mathbf{Q}_{n-1} + (1-\lambda)\max(\tilde{\mathbf{Q}}_{n-1}, \mathbf{Q}_n^{\min}) \tag{36}$$

Optimizing the stochastic forecast helps the assimilation adjust to changing ionospheric conditions, as all deviations from the background model are driven by the stochastic component of the forecast. This approach does have limitations, notably that an undirected random walk is just as likely to step a particle in the wrong direction as it is to move it closer to the truth. When the stochastic variance $\mathbf{Q}_n$ increases in response to disturbed conditions, the particle filter is attempting to sample a larger region of state space with a fixed number of samples. While this is necessary to model the behavior of the ionospheric system, it does result in reduced sampling efficiency as particles are given a greater ability to take random steps into unlikely states.

Figure 5. Diagram showing a simplified representation of the particle filtering technique used in A-CHAIM. The points follow the changes to the particles through a single assimilation step, with the diameter representing their weight. Panel (a) Show the particles from the previous assimilation step. In (b) the deterministic part of the forecast $\mathbf{F}_n(\mathbf{x}_n)$ brings the particles to $\mathbf{X}_{n|n-1}^i$. In (c) each particle is resampled multiple times, and the daughter particles $\breve{\mathbf{X}}_n^i$ are given preliminary weights $\breve{w}_n^{ij}$. The highest weight offspring of each original particle is selected to form $\mathbf{X}_n^i$, and in (d) the entire set of observations is used to generate the final weights w_n^i. A final resampling occurs in (e), to remove any low weight particles to pass $\tilde{\mathbf{X}}_n^i$ to the next assimilation step.

This problem can be alleviated if the number of samples of state space is not fixed. In A-CHAIM we use a subset of the available observations $\breve{\mathbf{y}}_n \subset \mathbf{y}_n$ with a greatly expanded number of particles to attempt to steer the particles into higher probability regions. It is computationally trivial to evaluate the measurement operator for ionosonde characteristics. Unlike the GNSS data, the ionosonde measurements are often direct measurements of an ionospheric profile parameter, which are linear functions of the A-CHAIM state.

After the deterministic component of the forecast step, A-CHAIM resamples uniformly from $\mathbf{X}_n^i$ to produce $\breve{N}$ temporary copies of each particle, hereafter labeled $\breve{\mathbf{X}}_n^{ij}$, $j = 1, 2, \ldots, \breve{N}$. In A-CHAIM, $\breve{N} = 20$ copies of each particle $\mathbf{X}^i$ are produced. Each of these daughter particles is then given the random displacement $\delta\breve{\mathbf{X}}_n^{ij} \sim \mathcal{N}(0, \mathbf{Q}_n)$. This gives a total of $N \cdot \breve{N} = 20000$ unique temporary particles.

For each temporary particle $\breve{\mathbf{X}}_n^{ij}$ a preliminary weight $\breve{w}^{ij}$ is produced using only ionosonde and altimeter data, along with $f(\mathbf{x}_n|\mathbf{x}_{n-1})$ and $q(\mathbf{x}_n|\mathbf{x}_{n-1})$. This allows A-CHAIM to predict which forecast steps are likely to produce particles with high weights. Using these weights a new set of particles is generated.

$$\left\{ \mathbf{X}_n^i \right\} = \left\{ \breve{\mathbf{X}}_n^{ik} | \breve{w}^{ik} = \sup \breve{w}^{ij}, \ j \in 1 \ldots \breve{N} \right\} \tag{37}$$

For each original particle $\mathbf{X}_n^i$, the highest weight daughter particle is kept, and all others discarded. This preliminary filtering is invisible to the rest of the particle filter, as preserves the one-to-one relationship between $\mathbf{X}_{n-1}^i$ and $\mathbf{X}_n^i$, and so the weights $W_{1:n-1}$ are unchanged. This is in contrast to the more conventional resampling in Equations 23 and 24, where the weights $W_{1:n}$ are reset to $\frac{1}{N}$. The $\mathbf{X}_n^i$ produced by this forecast sampling are indistinguishable from those which could be found by generating N random displacements from $\mathcal{N}(0, \mathbf{Q}_n)$, except that they are less likely to have arrived in a low weight region of state space by chance.

Figure 5 gives a schematic overview of how this process integrates with the rest of A-CHAIM. The forecast sampling takes less than a second of computation time per assimilation step, but allows the filter to behave as if it had a factor of $\breve{N}$ more particles. Most random displacements result in suboptimal particles that can be easily rejected by ionosonde data. This ensures that computationally expensive sTEC raytracing is not wasted.

4. Results

To evaluate the performance of the assimilation, we will compare the predicted values at each latency to data sources which were not included in the assimilation. The assimilation should produce an improved representation of ionospheric electron density. In particular, this assimilation should address some of the known shortcomings of the background model E-CHAIM that were outlined in the introduction, in particular improving the spatial and temporal resolution.

Table 3
Locations of Ionosondes Used to Validate A-CHAIM Performance

Station	Geographic coords.	AACGM coords. (300 km)	Source link
Pond Inlet (PONC)	72.69°N, 282.04°E	80.18°N, 1.87°E	chain.physics.unb.ca
Svalbard (SVAL)	78.93°N, 11.85°E	77.00°N, 106.58°E	www.tgo.uit.no
Sodankylä (SODAN)	67.4°N, 26.6°E	65.0°N, 105.9°E	www.sgo.fi
Blissville (BLISS)	45.61°N, 293.46°E	53.43°N, 14.45°E	chain.physics.unb.ca

4.1. Ionosondes

Ionosondes which do not provide automatically processed data, and therefore are not available for assimilation, provide an ideal reference to examine the performance of A-CHAIM. Table 3 summarizes the geographic and magnetic coordinates of the four reference stations used in this study, and they are also shown along with the assimilated data in Figure 1. The instruments in Pond Inlet and Blissville are Canadian Advanced Digital Ionosondes (CADIs) (Jayachandran et al., 2009). The ionosonde on Svalbard, located in Ny-Ålesund, is also a CADI. It is operated by Tromsø Geophysical Observatory. Ionograms from these instruments were manually processed at 30-min time resolution and subsequently inverted to extract hmF2 using the POLynomial ANalysis (POLAN) software package (Titheridge, 1988). In addition to these three CADI systems, we will also use the Alpha-Wolf ionosonde at Sodankylä, operated by the Sodankylä Geophysical Observatory (SGO). The instrument is well positioned to provide an assessment of system performance in the European sector, and manually processed hourly data is publicly available.

The value of foF2 measured at each station though the assimilation period are plotted in Figure 6. In order to highlight the differences between A-CHAIM and E-CHAIM, the right column shows the same data with the value predicted by E-CHAIM subtracted. This serves primarily to remove the diurnal variation in foF2, which is well captured by E-CHAIM. The stations are ordered by decreasing geomagnetic latitude, with PONC and SVAL in the polar cap, SODAN at auroral latitudes, and BLISS in the midlatitudes.
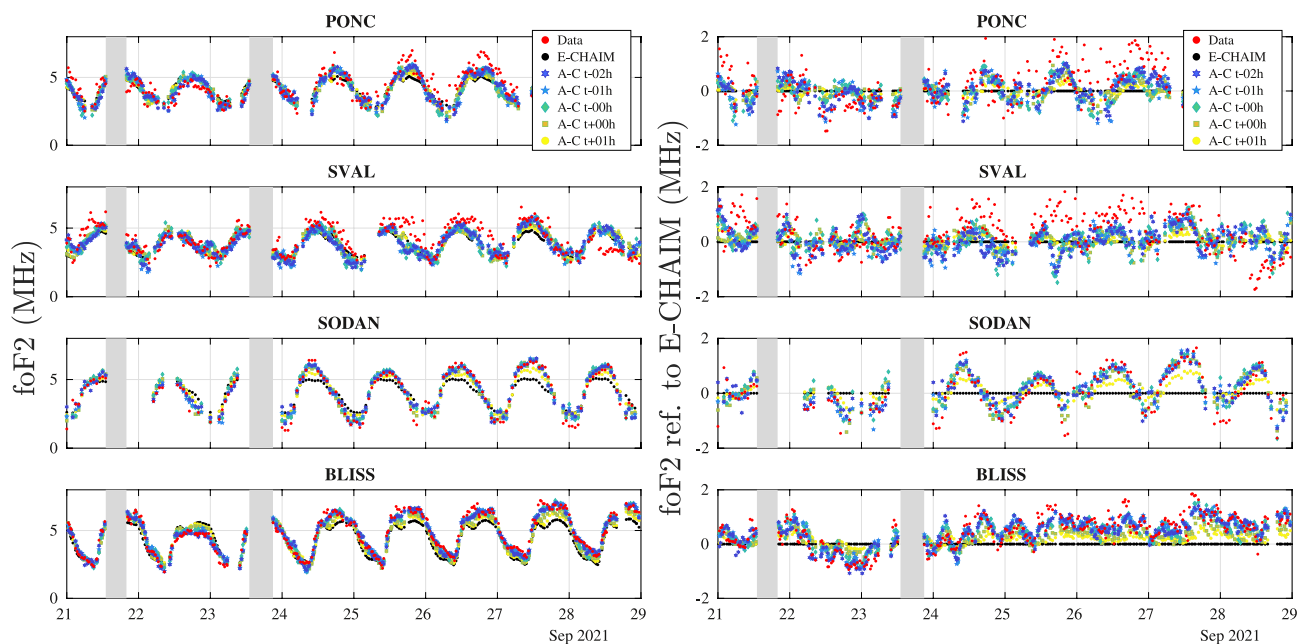

Figure 6. A-CHAIM foF2 performance at four reference ionosondes during the 21 September through 29 September 2021 period. The left column shows the predicted values of A-CHAIM and E-CHAIM plotted against the manually processed observations. The right column shows the same data with the value of E-CHAIM subtracted, to remove the diurnal variations. Gray bars mark periods where at least one of the latencies was not available due to missing data. For clarity, only the values of E-CHAIM t−02h are plotted.

Table 4
Summary of A-CHAIM and E-CHAIM Performance in foF2 Determination at Four Reference Ionosondes During the September 21st Through 29 September 2021 Period

Station	RMSE	t−02h	t−01h	t−00h	t+00h	t+01h
PONC	A-CHAIM	0.69	0.71	0.73	0.73	0.76
	E-CHAIM	0.80	0.80	0.79	0.79	0.79
	Δ (MHz)	−0.12	−0.09	−0.06	−0.06	−0.03
	Δ (%)	−14.4	−11.7	−7.4	−8.1	−3.8
SVAL	A-CHAIM	0.85	0.83	0.84	0.79	0.79
	E-CHAIM	0.83	0.83	0.80	0.80	0.80
	Δ (MHz)	0.02	0.00	0.04	−0.01	−0.01
	Δ (%)	2.1	0.6	4.6	−1.1	−1.8
SODAN	A-CHAIM	0.48	0.42	0.43	0.39	0.54
	E-CHAIM	0.75	0.75	0.76	0.76	0.76
	Δ (MHz)	−0.27	−0.33	−0.33	−0.36	−0.22
	Δ (%)	−36.1	−43.7	−43.0	−48.1	−28.7
BLISS	A-CHAIM	0.34	0.32	0.33	0.47	0.59
	E-CHAIM	0.80	0.78	0.75	0.75	0.76
	Δ (MHz)	−0.46	−0.46	−0.43	−0.28	−0.16
	Δ (%)	−57.2	−58.7	−56.6	−37.7	−21.8

Note. The rows summarize the overall RMSE in MHz at each latency for A-CHAIM and E-CHAIM, those labeled Δ show the difference in RMSE both in MHz, and as a percentage of the E-CHAIM RMSE.

The variability of the high-latitude ionosphere is immediately apparent at PONC and SVAL, with the variation between sequential measurements being much larger than the difference between A-CHAIM and E-CHAIM. As this data came from manually scaled ionograms, this variation is not an artifact of the autoscaling process, but is a result of the dynamic processes of the storm. Each of these stations is also in a relatively data-sparse region. The ionosonde THJ76 at Thule is relatively close to PONC, and several GNSS receivers are nearby, including a co-located receiver. Only a single GNSS receiver on Svalbard provided data near the SVAL ionosonde, which was not enough to provide a meaningful improvement. This can be seen in Table 4, where the overall RMSE of each latency of A-CHAIM and E-CHAIM are tabulated for each station. PONC shows a marginal overall improvement at each latency, with a reduction in error of 0.12 MHz at the t−02h latency, with worsening performance as the latency decreases. SVAL shows essentially no change in overall performance at any latency. While this does demonstrate a limitation of the assimilation, this is desirable behavior. If the assimilation is not able to improve upon the background model, either due to a lack of data, or because the state space we have chosen is not able to capture the real ionospheric behavior, then the ideal result would be to make no changes to the background.

The relative performance of A-CHAIM becomes very different once we move to lower geomagnetic latitudes. SODAN is situated near several ionosondes which were included in the assimilation, as well as the dense GNSS networks in Europe. While the variability in foF2 is lower than at the higher-latitude stations, we can see that E-CHAIM consistently underestimates foF2 during the day, and overestimates foF2 at night during this period. As a result, the overall E-CHAIM RMSE at SODAN is comparable to both PONC and SVAL at 0.8 MHz. At SODAN, A-CHAIM shows a strong improvement over E-CHAIM at all latencies, with RMSE between 0.39 and 0.54 MHz. Unexpectedly, the best performance at this location is at one of the forecasted latencies, t+00h. This effect is small and unique to this station, which has fewer total observations than the other reference stations in this study.

The improvement in foF2 that A-CHAIM produces is readily apparent at the mid-latitude BLISS ionosonde. E-CHAIM consistently overestimates the peak electron density, except during the depletion from September 22nd and 23rd, where E-CHAIM overestimates the peak. A-CHAIM is able to correct this diurnal-scale error. Additionally, A-CHAIM is able to capture smaller time-scale variations, most notably through September 24th. Overall A-CHAIM shows a strong improvement at all latencies, reducing the error from 0.75 to 0.3 MHz at latencies with available data. The forecasted latencies show diminishing improvements, with the most advanced forecast showing an error of 0.5 MHz.

We can also examine the ability of A-CHAIM to model the altitude of the peak electron density, hmF2. The results of the assimilation at each station are plotted in Figure 7. There are no observations of hmF2 at SODAN in this analysis, as the electron density profiles are not inverted by the SGO. The overall RMSE for the remaining three stations are given in Table 5, for each latency of A-CHAIM and E-CHAIM. Every station reporting data shows the same overall trend, namely a small but consistent improvement in hmF2 across all latencies, with the performance of each latency directly influenced by the amount of available data. While this is the behavior we would expect from a well-condition assimilation, these results are more striking when we compare them to the foF2 results in Table 4.

At SVAL there was no meaningful change in the foF2 RMSE, whereas hmF2 RMSE at that location was reduced by 12%–16%. GNSS sTEC and JASON measurements are not sensitive to changes in hmF2, and so most of the improvement in hmF2 must be driven by assimilated ionosonde measurements of hmF2. Given that our reference stations, and in particular SVAL, are isolated from other ionosondes, these improvements must be a result of large spatial scale corrections to hmF2. This is an advantage of fitting a parameter to a global basis set, improvements can be projected far from where the observations were made. A natural corollary to this advantage is that a global basis set also allows errors to be projected anywhere in the assimilation region. Investigating this possibility thoroughly requires using a reference data set with more global coverage than ionosondes can provide.
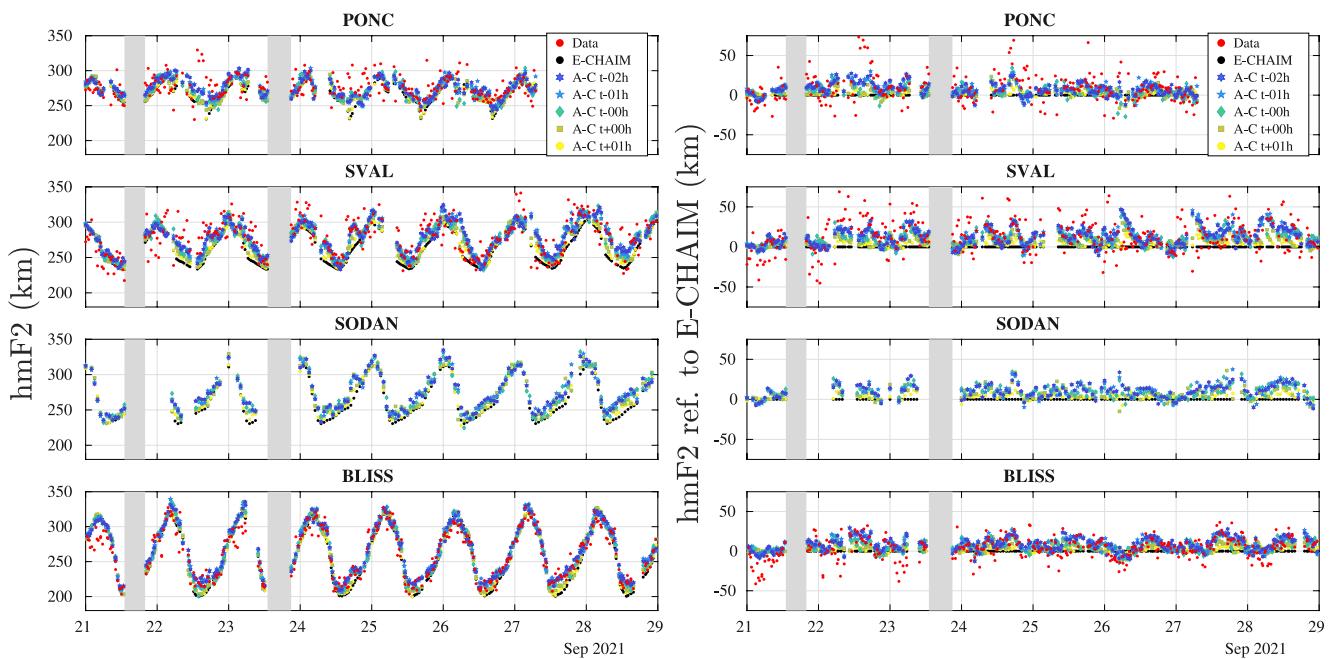
Figure 7. A-CHAIM hmF2 performance at four reference ionosondes during the September 21st through 29 September 2021 period. The left column shows the predicted values of A-CHAIM and E-CHAIM plotted against the manually processed observations. The right column shows the same data with the value of E-CHAIM subtracted, to remove the diurnal variations. Gray bars mark periods where at least one of the latencies was not available due to missing data. For clarity, only the values of E-CHAIM t−02h are plotted.

4.2. In Situ Electron Density Measurements

In order to assess the performance of the assimilation in the entire region, we will also make use of a limited data set of in situ plasma density measurements made onboard the Defense Meteorological Satellite Program (DMSP) and European Space Agency's Swarm satellite missions. These instruments have several properties that make them ideal as a rigorous test of performance, the foremost being their global coverage. This gives an ability to test the assimilation over the regions where no ground-based data is available, particularly over the oceans. As well, no measurements of in-situ plasma density, nor other direct measurements from the topside ionosphere, are included in the assimilated data. If the assimilation is altering the ionospheric state in an unphysical or inconsistent way, then these in-situ measurements would provide an ideal test.

4.2.1. Swarm

The ESA Swarm mission is a constellation of three satellites (Swarm A, B, C). Each satellite operates in a polar orbit with slow local time precession of 2.7 hr/month (Knudsen et al., 2017). As of 2021, Swarm A and Swarm C orbit at 440 km and Swarm B orbits at 505 km. As Swarm does not provide data in near real time, this study will make use of Swarm Langmuir Probe in situ measurements for independent validation of the assimilation system. To prepare the data, the Lomidze et al. (2018) calibration factors have been applied to the data set prior to comparison and all periods with non-nominal quality flags were discarded.

Figure 8 shows the root mean square error (RMSE) of all three Swarm satellites for all latencies of E-CHAIM and A-CHAIM, binned by geographic latitude and longitude for the entire study period. The errors in E-CHAIM are concentrated in three regions, the outer edge of the model where

Table 5
Summary of A-CHAIM and E-CHAIM Performance in hmF2 Determination at Three Reference Ionosondes During the 21 September Through 29 September 2021 Period

Station	RMSE	t−02h	t−01h	t−00h	t+00h	t+01h
PONC	A-CHAIM	18.43	18.91	19.47	19.49	19.60
	E-CHAIM	20.67	20.68	20.68	20.68	20.68
	Δ (km)	−2.24	−1.76	−1.21	−1.19	−1.08
	Δ (%)	−10.8	−8.5	−5.9	−5.8	−5.2
SVAL	A-CHAIM	21.29	21.41	21.31	21.04	22.12
	E-CHAIM	25.18	25.19	25.19	25.19	25.19
	Δ (km)	−3.90	−3.77	−3.87	−4.15	−3.07
	Δ (%)	−15.5	−15.0	−15.4	−16.5	−12.2
BLISS	A-CHAIM	12.84	13.32	13.49	13.85	14.03
	E-CHAIM	14.67	14.67	14.67	14.68	14.68
	Δ (km)	−1.83	−1.36	−1.18	−0.83	−0.65
	Δ (%)	−12.5	−9.2	−8.0	−5.6	−4.4

Note. The rows summarize the overall RMSE in km at each latency for A-CHAIM and E-CHAIM, those labeled Δ show the difference in RMSE both in km, and as a percentage of the E-CHAIM RMSE.
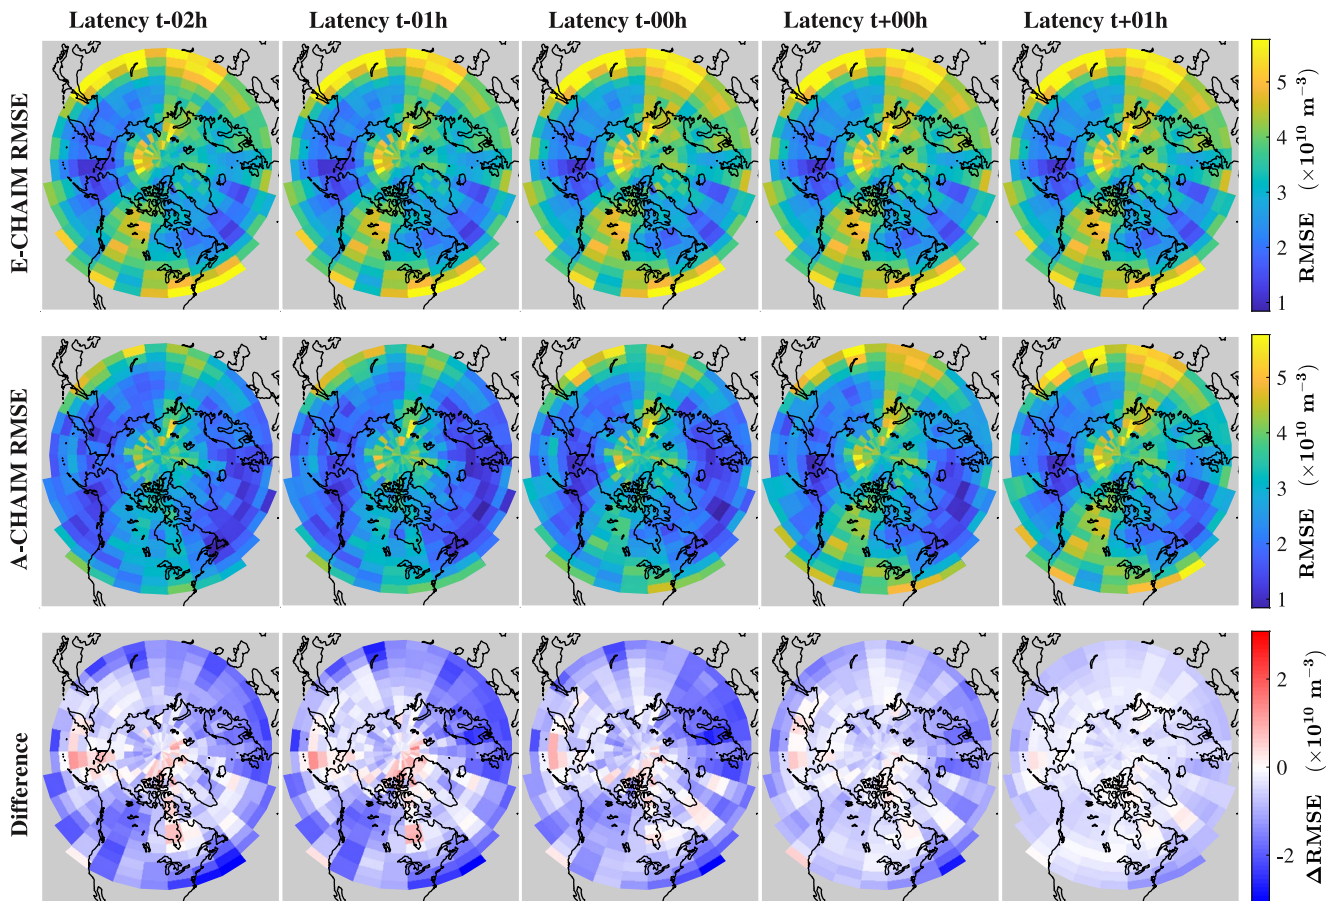
Figure 8. A-CHAIM performance using in-situ electron density measurements from the Swarm A, Swarm B and Swarm C satellites during the 21 September through 29 September 2021 period, binned by latitude and longitude. The top row shows the overall E-CHAIM RMSE, and the middle row shows the overall A-CHAIM RMSE. The bottom row shows E-CHAIM RMSE subtracted from the A-CHAIM RMSE at each latency, to highlight the differences. Measurements during periods where one or more of the latencies were unavailable were excluded.

$45° < MLAT < 50°$, the polar cap, and over central Canada. These patterns are consistent across all latencies, with slightly deteriorating performance at t−00h and later when the storm model was unavailable. By comparison, the errors in A-CHAIM are more spatially uniform at all latencies where data is available. The forecasted A-CHAIM (t+00h, t+01h) are still more uniform than the background model, but do show a steady decrease in performance relative to the assimilated latencies (t−02h, t−01h, t−00h).

The errors in A-CHAIM are significantly reduced compared to E-CHAIM at each latency. The greatest change is in the low magnetic latitude region, and over central Canada and Eurasia. Of note is the significant improvement at low latitudes over the Atlantic and Pacific Ocean. There are some regions which show a slight decrease in performance, namely over large bodies of water at high latitudes, for example, the Bering Sea, Hudson Bay, and along the northern coast of Greenland. These tend to be areas where there are relatively few measurements available, where E-CHAIM does relatively well, and nearby to regions where E-CHAIM does particularly poorly.

Table 6 summarizes the overall RMSE for each Swarm satellite, model, and latency. The strict ordering of the model performance is notable. For each satellite, every latency of A-CHAIM has a lower overall error than any version of E-CHAIM, with or without the storm model. The performance of A-CHAIM is always best at the t−02h latency, with decreasing performance as less data is available. The difference in RMSE between E-CHAIM with and without the storm model is also evident. The overall error is reduced by 15%–20% for all latencies other than the longest forecast t+01h, with a more modest 8%–9% reduction.

Table 6
Summary of A-CHAIM and E-CHAIM Performance Using In Situ Electron Density Measurements From the Swarm A, Swarm B and Swarm C Satellites During the 21 September Through 29 September 2021 Period

Satellite	RMSE	t−02h	t−01h	t−00h	t+00h	t+01h
Swarm A	A-CHAIM	3.06	3.12	3.24	3.43	3.64
	E-CHAIM	3.76	3.78	3.99	3.99	4.00
	$\Delta \left(m^{-3} \times 10^{10} \right)$	−0.70	−0.66	−0.76	−0.56	−0.36
	Δ (%)	−18.5	−17.4	−19.0	−14.1	−8.9
Swarm B	A-CHAIM	2.37	2.40	2.61	2.79	3.01
	E-CHAIM	3.05	3.06	3.29	3.29	3.29
	$\Delta \left(m^{-3} \times 10^{10} \right)$	−0.68	−0.66	−0.68	−0.50	−0.28
	Δ (%)	−22.2	−21.5	−20.8	−15.1	−8.5
Swarm C	A-CHAIM	2.98	3.04	3.10	3.26	3.45
	E-CHAIM	3.56	3.57	3.78	3.78	3.78
	$\Delta \left(m^{-3} \times 10^{10} \right)$	−0.58	−0.54	−0.68	−0.52	−0.33
	Δ (%)	−16.3	−15.1	−17.9	−13.7	−8.7

Note. The rows summarize the overall RMSE in $m^{-3} \times 10^{10}$ at each latency for A-CHAIM and E-CHAIM, those labeled Δ show the difference in RMSE both in absolute terms, and as a percentage of the E-CHAIM RMSE.

4.2.2. DMSP

The DMSP satellites (F-16, F-17, and F-18) orbit in a Sun-synchronous, circular orbit at between 830 and 880 km, each with an orbital period of 110 min (Garner et al., 2010). Similar to Swarm, the DMSP satellites each operate an array of in situ plasma density measurement systems and also do not provide data in near real time. Due to their higher altitude, these satellites represent a unique validation data set, and given the strong performance of the E-CHAIM background model in comparison to DMSP in the past (Themens, Jayachandran, McCaffrey, Reid, & Varney, 2019), this data set should pose a significant challenge to achieving improvement over the background.

Figure 9 follows the same format as Figure 8, showing the binned RMSE for all latencies of E-CHAIM and A-CHAIM, binned by geographic latitude and longitude for the entire study period. The performance of E-CHAIM varies strongly with latitude, with relatively minor variations longitudinally. The greatest errors are concentrated in a ring in the polar cap, and to a lesser degree at the outer edge of the assimilation region $MLAT < 50°$. A-CHAIM preserves this overall form, with the greatest errors at the extreme high and low latitudes. The errors in A-CHAIM are more evenly distributed across the assimilation region when compared to E-CHAIM, which is similar to the trend observed in the Swarm data.

The errors in A-CHAIM are significantly reduced at each latency, when compared to the corresponding E-CHAIM result. There are strong improvements at virtually all auroral and sub-auroral latitudes, including over the Atlantic Ocean, the Russian Far East and much of the Pacific. The greatest improvements occur in the American and European sectors. There are several regions that do show a slight decrease in performance, namely over the Pacific near North America, over the southern tip of Greenland, and in a few places over the Arctic ocean in the European and Russian sectors. As in the Swarm data, these are places where there are few measurements, where E-CHAIM performed relatively well, and are in close proximity to regions with comparatively large errors.

Table 7 summarizes the overall RMSE for each DMSP satellite, model, and latency. For each satellite, A-CHAIM t−02h has the best performance, and each latency that follows shows a decrease in performance as fewer observations are available. Every latency of A-CHAIM has smaller overall error than any of the E-CHAIM latencies, as we saw in the Swarm data. The overall error is reduced by 24%–29% for all latencies with assimilated data. The t+00h forecast has a reduction in error of 19% and the longest forecast t+01h shows a 10% reduction.

It is clear from examining the in-situ data from both DMSP and Swarm that A-CHAIM is able to provide a significant improvement in electron density at all latencies, including in regions where no observations are available, and during forecasts. In addition to an overall reduction in error, the spatial distribution of errors is more

Figure 9. A-CHAIM performance using in-situ electron density measurements from the DMSP F-16, F-17, and F-18 satellites during the September 21st through 29 September 2021 period, binned by latitude and longitude. The top row shows the overall E-CHAIM RMSE, and the middle row shows the overall A-CHAIM RMSE. The bottom row shows E-CHAIM RMSE subtracted from the A-CHAIM RMSE at each latency, to highlight the differences. Measurements during periods where one or more of the latencies were unavailable were excluded.

Table 7
Summary of A-CHAIM and E-CHAIM Performance Using In Situ Electron Density Measurements From the DMSP F-16, F-17, and F-18 Satellites During the 21 September Through 29 September 2021 Period

Satellite	RMSE	t−02h	t−01h	t−00h	t+00h	t+01h
F-16	A-CHAIM	5.57	5.75	6.43	7.09	7.83
	E-CHAIM	7.86	7.99	8.74	8.74	8.74
	$\Delta\ (m^{-3} \times 10^{9})$	−2.29	−2.24	−2.31	−1.65	−0.92
	$\Delta\ (\%)$	−29.2	−28.0	−26.4	−18.9	−10.5
F-17	A-CHAIM	5.79	5.86	6.40	6.87	7.55
	E-CHAIM	7.61	7.71	8.39	8.39	8.39
	$\Delta\ (m^{-3} \times 10^{9})$	−1.82	−1.86	−1.98	−1.52	−0.84
	$\Delta\ (\%)$	−23.9	−24.1	−23.7	−18.1	−10.0
F-18	A-CHAIM	5.23	5.44	6.05	6.56	7.29
	E-CHAIM	7.33	7.48	8.16	8.16	8.16
	$\Delta\ (m^{-3} \times 10^{9})$	−2.10	−2.03	−2.11	−1.61	−0.87
	$\Delta\ (\%)$	−28.7	−27.2	−25.8	−19.7	−10.7

Note. The rows summarize the overall RMSE in $m^{-3} \times 10^{9}$ at each latency for A-CHAIM and E-CHAIM, those labeled Δ show the difference in RMSE both in absolute terms, and as a percentage of the E-CHAIM RMSE.

even than in E-CHAIM. There were some areas where the performance of A-CHAIM was slightly worse than in E-CHAIM in at least one of the datasets. In both DMSP and Swarm, this only occurred in areas that had few observations, good E-CHAIM performance, and were in close proximity to regions with poor E-CHAIM performance. This is likely inevitable due to the limited horizontal resolution of the model. By correcting the region with poor performance, and without sufficient observations to constrain it, the assimilation can disrupt a nearby region where the background model does unusually well. Equivalently, if the assimilation smooths out the spatial variation of the errors in the background model, then regions where the background model performs well may end up worse off, even as the overall error is significantly reduced.

5. Conclusion

This study was performed using both data and outputs from the background model E-CHAIM that were produced in real time, in an operational environment. Using these, A-CHAIM is able to produce a significant improvement in modeled electron density when compared to the background model E-CHAIM. This reduction in error is largely uniform across the entire assimilation region, as measured by in-situ satellite-borne electron density measurements. The performance of A-CHAIM is best at higher latencies, up to 3 hours before the current time. However, A-CHAIM is able to produce an improved representation of electron density in near-real-time, with a 15%–25% reduction in error. A-CHAIM is also able to show improvements up to 2 hours in the future as a low-skill forecast, with a 15%–20% reduction in error in the first hour, and 8%–10% reduction in the second hour.

The ability of A-CHAIM to describe the shape of the ionosphere was also assessed, using four manually-processed ionosondes that were not included in the assimilated data. The critical frequency of the F2 layer, foF2, shows strong improvement at mid- and auroral latitudes, but does not show a significant improvement in the polar cap. At the lower latitude stations, A-CHAIM was able to produce an improvement of 0.3–0.46 MHz in near-real-time, and a 0.15–0.2 MHz improvement in the second hour of the forecast. A-CHAIM is also able to improve hmF2 at all latitudes, although the scale of the improvement is small (<5 km) when compared to natural ionospheric variability.

The challenges created by sparse data, limited computing resources, and unknown physical drivers are not unique to A-CHAIM, or the high latitude ionosphere. The unique flexibility of particle filtering as a data assimilation technique can be used to circumvent some of these issues, as the above results demonstrate. While particle filters do have limitations, in particular weight degeneracy, the techniques developed for A-CHAIM should be broadly applicable. Reducing the dimensionality of the measurements by building composite observables should produce a strong improvement when assimilating large numbers of low-information observations, and can be used in conjunction with localization techniques in systems that admit them.

As A-CHAIM continues to operate, further studies will need to be taken to assess the long term trends in performance. A-CHAIM does also produce estimates of the DCBs of the GNSS receivers it assimilates, and characterization of the accuracy and stability of those biases needs to be evaluated.

Appendix A: χ^2 Distribution of Log-Likelihood

Let S be some physical system, characterized by some set of parameters $\mathbf{x}$. All possible configurations of S define a state space $\mathbb{X}$, such that $\mathbf{x} \in \mathbb{X}$. The properties of S can only be inferred through some set of observables $\mathbf{y}$, which exist in a measurement space $\mathbb{Y}$ such that $\mathbf{y} \in \mathbb{Y}$. There is a measurement operator $\mathbf{H}(\mathbf{x}) : \mathbb{X} \to \mathbb{Y}$ which maps from the state space to the measurement space. $\mathbf{H}(\mathbf{x})$ is also described as the forward model, as it is a model which predicts which set of observations $\mathbf{y}_{obs}$ would be produced for a given state $\mathbf{x}$. For example, for ionospheric slant Total Electron Content, $\mathbf{H}(\mathbf{x})$ could take the form in Equation A1, a set of line integrals for each of m receiver-satellite pairs.

$$\mathbf{H}(\mathbf{x}) = \left(\int_{\vec{r}_{rcv,i}}^{\vec{r}^{sat,i}} N_e(\mathbf{x}, \vec{r}) d\vec{r} \right)^T \text{ for } i \in 1, 2, \ldots m \tag{A1}$$

In general, both the parameterization of the system, and the construction of the forward model $\mathbf{H}(\mathbf{x})$ will be imperfect. This, along with any error in the measuring instrument itself, contribute to the measurement error ϵ. If $\mathbf{x}_{true}$ is the true configuration of the system, then the actual values observed $\mathbf{y}_{obs}$ will be Equation A2.

$$\mathbf{y}_{obs} = \mathbf{H}(\mathbf{x}_{true}) + \epsilon \tag{A2}$$

If the measurement error is able to be represented by a multivariate Gaussian distribution, then ϵ is a vector of m random variables with $\epsilon \sim \mathcal{N}(0, \mathbf{R})$. This gives the familiar form Equation A3 for the likelihood $p(\mathbf{y}|\mathbf{x})$ of observations with multivariate Gaussian error.

$$p(\mathbf{y}|\mathbf{x}) = \exp\left[-\frac{1}{2}(\mathbf{y} - \mathbf{H}(\mathbf{x}))^T \mathbf{R}^{-1}(\mathbf{y} - \mathbf{H}(\mathbf{x}))\right] \tag{A3}$$

As $\mathbf{R}$ must be positive definite, it has a Cholesky decomposition $\mathbf{R} = \mathbf{L}\mathbf{L}^T$. If $\mathbf{z}$ is a vector of i.i.d. random variables $\mathbf{z} \sim \mathcal{N}(0, \mathbb{I}_m)$

$$\epsilon = \mathbf{y} - \mathbf{H}(\mathbf{x}) = \mathbf{L}\mathbf{z} \tag{A4}$$

$$\begin{aligned}
\log p(\mathbf{y}|\mathbf{x}) &= -\frac{1}{2}(\mathbf{y} - \mathbf{H}(\mathbf{x}))^T \mathbf{R}^{-1}(\mathbf{y} - \mathbf{H}(\mathbf{x})) \\
&= -\frac{1}{2}(\mathbf{L}\mathbf{z})^T\left(\mathbf{L}\mathbf{L}^T\right)^{-1}(\mathbf{L}\mathbf{z}) \\
&= -\frac{1}{2}\mathbf{z}^T\mathbf{L}^T\left(\mathbf{L}^T\right)^{-1}\mathbf{L}^{-1}\mathbf{L}\mathbf{z} \\
&= -\frac{1}{2}\mathbf{z}^T\mathbf{z}
\end{aligned} \tag{A5}$$

Let $l(\mathbf{x}) = -2\log p(\mathbf{y}|\mathbf{x})$, then $l(\mathbf{x})$ is proportional to a sum of squared standard normal variables, meaning it is χ^2 distributed with $n(\mathbf{z}) = m$ degrees of freedom.

Data Availability Statement

The near real time outputs of A-CHAIM, along with software to interpret the output files, is publicly available at https://www.rspl.ca/index.php/projects/chaim/a-chaim. Interpreter software is available in the C and MATLAB languages. E-CHAIM is available at https://www.rspl.ca/index.php/projects/chaim/e-chaim, and is available in C, MATLAB, and IDL. The output files, interpreter, and all reference datasets used in this work are available at https://doi.org/10.5281/zenodo.6642849. The GNSS data used in A-CHAIM is provided by: the German Federal Agency for Cartography and Geodesy (BKG) for the International GNSS Service (IGS) (2021) https://igs.bkg.bund.de/root_ftp/IGS/highrate/, IAG (International Association of Geodesy) Regional Reference Frame sub-commission for Europe (EUREF) (2021) https://igs.bkg.bund.de/root_ftp/EUREF/highrate/, and Integrated Geodetic Reference Network of Germany (GREF) (2021) https://igs.bkg.bund.de/root_ftp/GREF/nrt/ networks; the Canadian High Arctic Ionospheric Network (CHAIN) (2021) http://chain.physics.unb.ca/data/gps/data/highrate/; the Crustal Dynamics Data Information System (CDDIS) (2021) https://cddis.nasa.gov/archive/gnss/data/highrate/; the NOAA National Geodetic Survey (NGS) (2021) http://geodesy.noaa.gov/corsdata/rinex/; the California Spatial Reference Center (CSRC) GARNER GPS Archive (2021) http://garner.ucsd.edu/pub/nrtdata/; Natural Resources Canada (NRCan) (2021) http://rtopsdata1.geod.nrcan.gc.ca/gps/data/nrtdata/; and the Ministry of Energy and Natural Resources (MERN) (2021) ftp://ftp.mrn.gouv.qc.ca/Public/GPS/. Precise orbit determination in.SP3 format is provided by International GNSS Service (IGS) (1994) https://cddis.nasa.gov/archive/gnss/products. Satellite DCBs are provided by the Institute of Geodesy and Geophysics (IGG) of the Chinese Academy of Sciences (CAS), International GNSS Service (IGS) (2013) https://cddis.nasa.gov/archive/gnss/products/bias/. Near-Real-Time Ionosonde data is provided by the National Centers for Environmental Information (NCEI) (2021b) https://www.ngdc.noaa.gov/ionosonde/data/; and by the Global Ionospheric Radio Observatory (GIRO) (2011) http://spase.info/SMWG/Observatory/GIRO. Altimeter data from the Jason-3 satellite is provided by the NOAA National Oceanographic Data Center https://www.ncei.noaa.gov/archive/accession/Jason3-xGDR. The CADI ionosonde data used for verification was provided by Vertical Incidence Soundings (Ionograms) (2021) http://chain.physics.unb.ca/data/cadi/, and the Tromsø Geophysical Observatory (2021) https://www.tgo.uit.no/ionosondeNAL/. Ionosonde data from Sodanklya was provided by the Sodankylä Geophysical Observatory (SGO) (2021) http://www.sgo.fi/pub_ion/dailydata/. In situ electron density measurements from the Swarm mission are provided by the European Space Agency (2021) at https://swarm-diss.eo.esa.int/#swarm%2FLevel1b%2FEntire_mission_data%2FEFIx_LP. In-situ measurements from the DMSP missions are provided by National Centers for Environmental Information (NCEI) (2021a) at https://satdat.ngdc.noaa.gov/dmsp/data/.

References

Altadill, D., Torta, J., & Blanch, E. (2009). Proposal of new models of the bottom-side $B0$ and $B1$ parameters for IRI. *Advances in Space Research*, *43*(11), 1825–1834. (Ionosphere - Modelling, Forecasting, and Telecommunications I). https://doi.org/10.1016/j.asr.2008.08.014

Angling, M. J., Shaw, J., Shukla, A. K., & Cannon, P. S. (2009). Development of an HF selection tool based on the electron density assimilative model near-real-time ionosphere. *Radio Science*, *44*(1), RS0A13. https://doi.org/10.1029/2008RS004022

Bengtsson, T., Bickel, P., & Li, B. (2008). Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems. *IMS Collections*, *2*, 316–334. https://doi.org/10.1214/193940307000000518

Bjoland, L. M., Belyey, V., Løvhaug, U. P., & La Hoz, C. (2016). An evaluation of international reference ionosphere electron density in the polar cap and cusp using EISCAT Svalbard radar measurements. *Annales Geophysicae*, *34*(9), 751–758. https://doi.org/10.5194/angeo-34-751-2016

Buchert, S. C. (2020). Entangled dynamos and joule heating in the Earth's ionosphere. *Annales Geophysicae*, *38*(5), 1019–1030. https://doi.org/10.5194/angeo-38-1019-2020

Bust, G. S., Garner, T. W., & Gaussiran, T. L., II. (2004). Ionospheric data assimilation three-dimensional (IDA3D): A global, multisensor, electron density specification algorithm. *Journal of Geophysical Research*, *109*(A11), A11312. https://doi.org/10.1029/2003JA010234

Bust, G. S., & Immel, T. J. (2020). Ida4d: Ionospheric data assimilation for the icon mission. *Space Science Reviews*, *216*(3), 33. https://doi.org/10.1007/s11214-020-00648-z

California Spatial Reference Center (CSRC) GARNER GPS Archive. (2021). GPS Rinex observation data in 30 sec hourly files [Dataset]. Garner. Retrieved from ftp://garner.ucsd.edu/pub/nrtdata/

Canadian High Arctic Ionospheric Network (CHAIN). (2021). GPS Rinex observation data in 1 sec hourly files [Dataset]. CHAIN. Retrieved from http://chain.physics.unb.ca/data/gps/data/highrate/

Carrano, C., & Groves, K. (2009). *Ionospheric data processing and analysis. Workshop on satellite navigation science and technology for Africa*. The Abdus Salam International Center for Theoretical Physics (ICTP).

Cervera, M. A., Francis, D. B., & Frazer, G. J. (2018). Climatological model of over-the-horizon radar. *Radio Science*, *53*(9), 988–1001. https://doi.org/10.1029/2018RS006607

Cervera, M. A., & Harris, T. J. (2014). Modeling ionospheric disturbance features in quasi-vertically incident ionograms using 3-D magnetoionic ray tracing and atmospheric gravity waves. *Journal of Geophysical Research: Space Physics*, *119*(1), 431–440. https://doi.org/10.1002/2013JA019247

Chartier, A. T., Kinrade, J., Mitchell, C. N., Rose, J. A. R., Jackson, D. R., Cilliers, P., et al. (2014). Ionospheric imaging in Africa. *Radio Science*, *49*(1), 19–27. https://doi.org/10.1002/2013RS005238

Chartier, A. T., Matsuo, T., Anderson, J. L., Collins, N., Hoar, T. J., Lu, G., et al. (2016). Ionospheric data assimilation and forecasting during storms. *Journal of Geophysical Research: Space Physics*, *121*(1), 764–778. https://doi.org/10.1002/2014JA020799

Cosgrove, R. B., & Codrescu, M. (2009). Electric field variability and model uncertainty: A classification of source terms in estimating the squared electric field from an electric field model. *Journal of Geophysical Research*, *114*(A6), A06301. https://doi.org/10.1029/2008JA013792

Coster, A., Williams, J., Weatherwax, A., Rideout, W., & Herne, D. (2013). Accuracy of GPS total electron content: GPS receiver bias temperature dependence. *Radio Science*, *48*(2), 190–196. https://doi.org/10.1002/rds.20011

Crustal Dynamics Data Information System (CDDIS). (2021). GPS Rinex observation data in 1 sec hourly files [Dataset]. CDDIS. Retrieved from https://cddis.nasa.gov/archive/gnss/data/highrate/

Douc, R., Cappé, O., & Moulines, E. (2005). Comparison of resampling schemes for particle filtering. CoRR, abs/cs/0507025. Retrieved from http://arxiv.org/abs/cs/0507025

Doucet, A., Godsill, S., & Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, *10*(3), 197–208. https://doi.org/10.1023/A:1008935410038

Doucet, A., & Johansen, A. (2009). A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of Nonlinear Filtering*, *12*.

Elvidge, S., & Angling, M. J. (2019). Using the local ensemble transform Kalman filter for upper atmospheric modelling. *Journal of Space Weather and Space Climate*, *9*(A30). https://doi.org/10.1051/swsc/2019018

European Space Agency. (2021). Langimur probe (LP) of the electrical field instrument (EFI) EFIX_LP_1B plasma product from swarm a, swarm b, and swarm c satellites [Dataset]. SWARM. Retrieved from https://swarm-diss.eo.esa.int/#swarm%2FLevel1b%2FEntire_mission_data%2FEFIx_LP

Fernandez-Gomez, I., Fedrizzi, M., Codrescu, M. V., Borries, C., Fillion, M., & Fuller-Rowell, T. J. (2019). On the difference between real-time and research simulations with CTIPE. *Advances in Space Research*, *64*(10), 2077–2087. (Variability and Coupling of the Equatorial, Low- and Mid-Latitude Mesosphere, Thermosphere and Ionosphere: Latest Developments of Monitoring and Modeling Techniques). https://doi.org/10.1016/j.asr.2019.02.028

Galkin, I. A., Reinisch, B. W., Huang, X., & Bilitza, D. (2012). Assimilation of GIRO data into a real-time IRI. *Radio Science*, *47*(4). https://doi.org/10.1029/2011RS004952

Garner, T. W., Taylor, B. T., Gaussiran, T. L., II, Coley, W. R., Hairston, M. R., & Rich, F. J. (2010). Statistical behavior of the topside electron density as determined from DMSP observations: A probabilistic climatology. *Journal of Geophysical Research*, *115*(A7), A07306. https://doi.org/10.1029/2009JA014695

Global Ionospheric Radio Observatory (GIRO). (2011). Giro tabulated ionospheric characteristics, version 1.0 revision b [Dataset]. SPACE. Retrieved from http://spase.info/SMWG/Observatory/GIRO

Hernández-Pajares, M., Juan, J. M., Sanz, J., Orus, R., Garcia-Rigo, A., Feltens, J., et al. (2009). The IGS VTEC maps: A reliable source of ionospheric information since 1998. *Journal of Geodesy*, *83*(3), 263–275. https://doi.org/10.1007/s00190-008-0266-1

Huang, X., & Reinisch, B. (2001). Vertical electron content from ionograms in real time. *Radio Science*, *36*(2), 335–342. https://doi.org/10.1029/1999RS002409

IAG (International Association of Geodesy) Regional Reference Frame sub-commission for Europe (EUREF). (2021). GPS Rinex observation data in 1 sec hourly files [Dataset]. BKG GNSS Data Center. Retrieved from https://igs.bkg.bund.de/root_ftp/EUREF/highrate/

Integrated Geodetic Reference Network of Germany (GREF). (2021). GPS Rinex observation data in 30 sec hourly files [Dataset]. BKG GNSS Data Center. Retrieved from https://igs.bkg.bund.de/root_ftp/GREF/nrt/

International GNSS Service (IGS). (1994). GNSS combined rapid solution orbit product [Dataset]. NASA Crustal Dynamics Data Information System. https://doi.org/10.5067/GNSS/GNSS_IGSRORB_001

International GNSS Service (IGS). (2013). GNSS differential code bias product [Dataset]. NASA Crustal Dynamics Data Information System. https://doi.org/10.5067/GNSS/GNSS_IGSDCB_001

International GNSS Service (IGS). (2021). GPS Rinex observation data in 1 sec hourly files [Dataset]. BKG GNSS Data Center. Retrieved from https://igs.bkg.bund.de/root_ftp/IGS/highrate/

Jayachandran, P. T., Langley, R. B., MacDougall, J. W., Mushini, S. C., Pokhotelov, D., Hamza, A. M., et al. (2009). Canadian high arctic ionospheric network (CHAIN). *Radio Science*, *44*(1), RS0A03. https://doi.org/10.1029/2008rs004046

Jones, W. B., & Gallet, R. M. (1962). Representation of diurnal and geographic variations of ionospheric data by numerical methods. *Journal of Research of the National Bureau of Standards*, *66*(D4), 419–438. https://doi.org/10.6028/jres.066d.043

Knudsen, D. J., Burchill, J. K., Buchert, S. C., Eriksson, A. I., Gill, R., Wahlund, J.-E., et al. (2017). Thermal ion imagers and Langmuir probes in the swarm electric field instruments. *Journal of Geophysical Research: Space Physics*, *122*(2), 2655–2673. https://doi.org/10.1002/2016JA022571

Lee, I. T., Matsuo, T., Richmond, A. D., Liu, J. Y., Wang, W., Lin, C. H., et al. (2012). Assimilation of FORMOSAT-3/COSMIC electron density profiles into a coupled thermosphere/ionosphere model using ensemble Kalman filtering. *Journal of Geophysical Research*, *117*(A10), A10318. https://doi.org/10.1029/2012JA017700

Li, M., Yuan, Y., Wang, N., Li, Z., & Huo, X. (2018). Performance of various predicted GNSS global ionospheric maps relative to GPS and Jason Tec data. *GPS Solutions*, *22*(2), 55. https://doi.org/10.1007/s10291-018-0721-2

Lockwood, M., Cowley, S. W. H., & Freeman, M. P. (1990). The excitation of plasma convection in the high-latitude ionosphere. *Journal of Geophysical Research*, *95*(A6), 7961–7972. https://doi.org/10.1029/JA095iA06p07961

Lomidze, L., Knudsen, D. J., Burchill, J., Kouznetsov, A., & Buchert, S. C. (2018). Calibration and validation of swarm plasma densities and electron temperatures using ground-based radars and satellite radio occultation measurements. *Radio Science*, *53*(1), 15–36. https://doi.org/10.1002/2017RS006415

Maltseva, O., & Nikitenko, T. (2021). Validation of various ionospheric models in the high-latitudinal zone. *Advances in Space Research*, *68*(5), 2233–2243. (International Reference Ionosphere - Progress and New Inputs). https://doi.org/10.1016/j.asr.2020.09.016

Ministry of Energy and Natural Resources (MERN). (2021). GPS Rinex observation data in 30 sec hourly files [Dataset]. FTP. Retrieved from http://mrn.gouv.qc.ca/Public/GPS/

Montenbruck, O., & González Rodríguez, B. (2019). NeQuick-G performance assessment for space applications. *GPS Solutions*, *24*(1), 13. https://doi.org/10.1007/s10291-019-0931-2

National Centers for Environmental Information (NCEI). (2021a). Bulk plasma parameters from the topside ionopheric plasma monitor (SSIES) instruments on the DMSP F-16, DMSP F-17 AND DMSP F-18 satellites [Dataset]. SATDAT. Retrieved from https://satdat.ngdc.noaa.gov/dmsp/data/

National Centers for Environmental Information (NCEI). (2021b). Vertical incidence soundings (ionograms). [Dataset]. National Oceanic and Atmospheric Administration. Retrieved from https://www.ngdc.noaa.gov/ionosonde/data/

Natural Resources Canada (NRCan). (2021). GPS Rinex observation data in 30 sec hourly files [Dataset]. NRCAN. Retrieved from http://rtopsdata1.geod.nrcan.gc.ca/gps/data/nrtdata/

Nesterov, I., & Kunitsyn, V. (2011). GNSS radio tomography of the ionosphere: The problem with essentially incomplete data. *Advances in Space Research*, *47*(10), 1789–1803. https://doi.org/10.1016/j.asr.2010.11.034

Nickisch, L. J., Fridman, S., Hausman, M., Kraut, S., & Zunich, G. (2016). Assimilative modeling of ionospheric dynamics for nowcasting of HF propagation channels in the presence of TIDS. *Radio Science*, *51*(3), 184–193. https://doi.org/10.1002/2015RS005902

NOAA National Geodetic Survey (NGS). (2021). GPS Rinex observation data in 30 sec hourly files. [Dataset]. National Oceanic and Atmospheric Administration. Retrieved from http://geodesy.noaa.gov/corsdata/rinex/

Ott, E., Hunt, B., Szunyogh, I., Zimin, A., Kostelich, E., Corazza, M., et al. (2004). A local ensemble Kalman filter for atmospheric data assimilation. *Tellus*, *56*(5), 415–428. https://doi.org/10.1111/j.1600-0870.2004.00076.x

Pignalberi, A., Pietrella, M., & Pezzopane, M. (2021). Towards a real-time description of the ionosphere: A comparison between international reference ionosphere (IRI) and IRI real-time assimilative mapping (IRTAM) models. *Atmosphere*, *12*(8), 1003. https://doi.org/10.3390/atmos12081003

Prol, F. S., Kodikara, T., Hoque, M. M., & Borries, C. (2021). Global-scale ionospheric tomography during the march 17, 2015 geomagnetic storm. *Space Weather*, *19*(12), e2021SW002889. https://doi.org/10.1029/2021SW002889

Rasmussen, C. E., Schunk, R. W., & Sojka, J. J. (1986). Effects of different convection models upon the high-latitude ionosphere. *Journal of Geophysical Research*, *91*(A6), 6999–7005. https://doi.org/10.1029/JA091iA06p06999

Reinisch, B. W., & Galkin, I. A. (2011). Global ionospheric radio observatory (GIRO). *Earth Planets and Space*, *63*(4), 377–381. https://doi.org/10.5047/eps.2011.03.001

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*(1). https://doi.org/10.1029/2002RS002797

Schunk, R. W., Scherliess, L., Eccles, V., Gardner, L. C., Sojka, J. J., Zhu, L., et al. (2016). Space weather forecasting with a multimodel ensemble prediction system (MEPS). *Radio Science*, *51*(7), 1157–1165. https://doi.org/10.1002/2015RS005888

Shepherd, S. G. (2014). Altitude-adjusted corrected geomagnetic coordinates: Definition and functional approximations. *Journal of Geophysical Research: Space Physics*, *119*(9), 7501–7521. https://doi.org/10.1002/2014JA020264

Shim, J. S., Kuznetsova, M., Rastätter, L., Hesse, M., Bilitza, D., Butala, M., et al. (2011). Cedar electrodynamics thermosphere ionosphere (ETI) challenge for systematic assessment of ionosphere/thermosphere models: Nmf2, hmf2, and vertical drift using ground-based observations. *Space Weather*, *9*(12), S12003. https://doi.org/10.1029/2011SW000727

Shim, J. S., Tsagouri, I., Goncharenko, L., Rastaetter, L., Kuznetsova, M., Bilitza, D., et al. (2018). Validation of ionospheric specifications during geomagnetic storms: Tec and foF2 during the 2013 march storm event. *Space Weather*, *16*(11), 1686–1701. https://doi.org/10.1029/2018SW002034

Smith, D. A., Araujo-Pradere, E. A., Minter, C., & Fuller-Rowell, T. (2008). A comprehensive evaluation of the errors inherent in the use of a two-dimensional shell for modeling the ionosphere. *Radio Science*, *43*(6), 1–23. https://doi.org/10.1029/2007RS003769

Sodankylä Geophysical Observatory (SGO). (2021). Vertical incidence soundings (ionograms) [Dataset]. SGO. Retrieved from http://www.sgo.fi/pub_ion/dailydata/

Spencer, P. S. J., & Mitchell, C. N. (2003). A three-dimensional time-dependent algorithm for ionospheric imaging using GPS (Vol. *46*(4)). https://doi.org/10.4401/ag-4373

Thayaparan, T., Ibrahim, Y., Polak, J., & Riddolls, R. (2018). High-frequency over-the-horizon radar in Canada. *IEEE Geoscience and Remote Sensing Letters*, *15*(11), 1700–1704. https://doi.org/10.1109/LGRS.2018.2856185

Themens, D. R., Jayachandran, P., Galkin, I., & Hall, C. (2017). The empirical Canadian high arctic ionospheric model (E-CHAIM): NMF2 and HMF2. *Journal of Geophysical Research: Space Physics*, *122*(8), 9015–9031. https://doi.org/10.1002/2017ja024398

Themens, D. R., Jayachandran, P., & Langley, R. (2015). The nature of GPS differential receiver bias variability: An examination in the polar cap region. *Journal of Geophysical Research: Space Physics*, *120*(9), 8155–8175. https://doi.org/10.1002/2015JA021639

Themens, D. R., Jayachandran, P., McCaffrey, A. M., Reid, B., & Varney, R. H. (2019). A bottomside parameterization for the empirical Canadian high arctic ionospheric model. *Radio Science*, *54*(5), 397–414. https://doi.org/10.1029/2018RS006748

Themens, D. R., Jayachandran, P., Nicolls, M., & MacDougall, J. (2014). A top to bottom evaluation of IRI-2007 within the polar cap. *Journal of Geophysical Research: Space Physics*, *119*(8), 6689–6703. https://doi.org/10.1002/2014JA020052

Themens, D. R., Jayachandran, P., Reid, B., & McCaffrey, A. (2020). The limits of empirical electron density modeling: Examining the capacity of E-CHAIM and the IRI for modeling intermediate (1- to 30-day) timescales at high latitudes. *Radio Science*, *55*(4), e2018RS006763. https://doi.org/10.1029/2018RS006763

Themens, D. R., Jayachandran, P., & Varney, R. H. (2018). Examining the use of the NeQuick bottomside and topside parameterizations at high latitudes. *Advances in Space Research*, *61*(1), 287–294. https://doi.org/10.1016/j.asr.2017.09.037

Themens, D. R., & Jayachandran, P. T. (2016). Solar activity variability in the IRI at high latitudes: Comparisons with GPS total electron content. *Journal of Geophysical Research: Space Physics*, *121*(4), 3793–3807. https://doi.org/10.1002/2016JA022664

Themens, D. R., Jayachandran, P. T., Bilitza, D., Erickson, P. J., Häggström, I., Lyashenko, M. V., et al. (2018). Topside electron density representations for middle and high latitudes: A topside parameterization for E-CHAIM based on the NeQuick. *Journal of Geophysical Research: Space Physics*, *123*(2), 1603–1617. https://doi.org/10.1002/2017JA024817

Themens, D. R., Jayachandran, P. T., Langley, R. B., Macdougall, J. W., & Nicolls, M. J. (2013). Determining receiver biases in GPS-derived total electron content in the auroral oval and polar cap region using ionosonde measurements. *GPS Solutions*, *17*(3), 357–369. https://doi.org/10.1007/s10291-012-0284-6

Themens, D. R., Jayachandran, P. T., & McCaffrey, A. M. (2019). Validating the performance of the empirical Canadian high arctic ionospheric model (E-CHAIM) with in situ observations from DMSP and champ. *Journal of Space Weather and Space Climate*, *9*(A21), e2018RS006763. https://doi.org/10.1051/swsc/2019021

Themens, D. R., Reid, B., & Elvidge, S. (2022). Artist ionogram autoscaling confidence scores: Best practices. *Radio Science Letters*. https://doi.org/10.46620/22-0001

Themens, D. R., Reid, B., Jayachandran, P. T., Larson, B., Koustov, A. V., Elvidge, S., et al. (2021). E-CHAIM as a model of total electron content: Performance and diagnostics. *Space Weather*, *19*(11), e2021SW002872. https://doi.org/10.1029/2021SW002872

Titheridge, J. E. (1988). The real height analysis of ionograms: A generalized formulation. *Radio Science*, *23*(05), 831–849. https://doi.org/10.1029/RS023i005p00831

Tromsø Geophysical Observatory. (2021). Vertical incidence soundings (ionograms) [Dataset]. TGO. Retrieved from https://www.tgo.uit.no/ionosondeNAL/

van Leeuwen, P. J. (2019). Non-local observations and information transfer in data assimilation. *Frontiers in Applied Mathematics and Statistics*, *5*. https://doi.org/10.3389/fams.2019.00048

van Leeuwen, P. J., Künsch, H. R., Nerger, L., Potthast, R., & Reich, S. (2019). Particle filters for high-dimensional geoscience applications: A review. *Quarterly Journal of the Royal Meteorological Society*, *145*(723), 2335–2365. https://doi.org/10.1002/qj.3551

Vertical Incidence Soundings (Ionograms). (2021). Vertical incidence soundings (ionograms) [Dataset]. Canadian High Arctic Ionospheric Network (CHAIN). Retrieved from http://chain.physics.unb.ca/data/cadi/

Wang, N., Yuan, Y., Li, Z., Montenbruck, O., & Tan, B. (2016). Determination of differential code biases with multi-GNSS observations. *Journal of Geodesy*, *90*(3), 209–228. https://doi.org/10.1007/s00190-015-0867-4

Watson, C., Themens, D. R., & Jayachandran, P. T. (2021). Development and validation of precipitation enhanced densities for the empirical Canadian high arctic ionospheric model. *Space Weather*, *19*(10), e2021SW002779. https://doi.org/10.1029/2021SW002779

Wielgosz, P., Milanowska, B., Krypiak-Gregorczyk, A., & Jarmołowski, W. (2021). Validation of GNSS-derived global ionosphere maps for different solar activity levels: Case studies for years 2014 and 2018. *GPS Solutions*, *25*(3), 103. https://doi.org/10.1007/s10291-021-01142-x