



# IONONEST—A Bayesian approach to modeling the lower ionosphere

**DOI:**

[10.1002/2016RS005965](https://doi.org/10.1002/2016RS005965)

**Document Version**

Final published version

[Link to publication record in Manchester Research Explorer](#)

**Citation for published version (APA):**

Martin, P., Scaife, A. M. M., McKay, D., & McCrea, I. (2016). IONONEST—A Bayesian approach to modeling the lower ionosphere. *Radio Science*, 51(8), 1332-1349. <https://doi.org/10.1002/2016RS005965>

**Published in:**

Radio Science

**Citing this paper**

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

**General rights**

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

**Takedown policy**

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact [uml.scholarlycommunications@manchester.ac.uk](mailto:uml.scholarlycommunications@manchester.ac.uk) providing relevant details, so we can investigate your claim.





## RESEARCH ARTICLE

10.1002/2016RS005965

## Key Points:

- We present IONONEST, a Bayesian method for producing  $D$  region electron density height profiles from multifrequency absorption measurements
- This method can quantitatively distinguish between electron density height profile models
- We validate the method through comparison with EISCAT VHF data

## Correspondence to:

P. L. Martin,  
plm1g08@soton.ac.uk

## Citation:

Martin, P. L., A. M. M. Scaife, D. McKay, and I. McCrea (2016), IONONEST—A Bayesian approach to modeling the lower ionosphere, *Radio Sci.*, 51, 1332–1349, doi:10.1002/2016RS005965.

Received 29 JAN 2016

Accepted 14 JUL 2016

Accepted article online 21 JUL 2016

Published online 11 AUG 2016

## IONONEST—A Bayesian approach to modeling the lower ionosphere

Poppy L. Martin<sup>1,2</sup>, Anna M. M. Scaife<sup>2</sup>, Derek McKay<sup>3,4</sup>, and Ian McCrea<sup>4</sup>

<sup>1</sup>Department of Physics and Astronomy, University of Southampton, Southampton, UK, <sup>2</sup>Jodrell Bank Centre for Astrophysics, University of Manchester, Manchester, UK, <sup>3</sup>Sodankylä Geophysical Observatory, University of Oulu, Tähteläntie 62, Sodankylä, Finland, <sup>4</sup>STFC Rutherford Appleton Laboratory, Harwell Science and Innovation Campus, Didcot, UK

**Abstract** Obtaining high-resolution electron density height profiles for the  $D$  region of the ionosphere as a well-sampled function of time is difficult for most methods of ionospheric measurement. Here we present a new method of using multifrequency riometry data for producing  $D$  region height profiles via inverse methods. To obtain these profiles, we use the nested sampling technique, implemented through our code, IONONEST. We demonstrate this approach using new data from the Kilpisjärvi Atmospheric Imaging Receiver Array (KAIRA) instrument and consider two electron density models. We compare the recovered height profiles from the KAIRA data with those from incoherent scatter radar using data from the European Incoherent Scatter Facility (EISCAT) instrument and find that there is good agreement between the two techniques, allowing for instrumental differences.

## 1. Introduction


The ionosphere is an upper layer of the Earth's atmosphere that ranges from altitudes of around 60 km to 1000 km [Hargreaves, 1992] and is split into layers, known as the  $D$ ,  $E$ , and  $F$  regions. The  $D$  region is the least well understood of these layers. It is an area of the ionosphere with complex chemistry that extends from altitudes of 60 km to 95 km [Hunsucker and Hargreaves, 2002]. As the ionization in this layer is primarily caused by solar activity, the electron density in this region is very low during the night, though it increases quickly with height. The electron density in the  $D$  region is substantially less than the other layers of the ionosphere, and so it cannot be measured effectively via the most common methods for observing the ionosphere, such as the ionosonde and GPS networks [Friedrich and Rapp, 2009].

The primary method to measure the  $D$  region is by using rocket-borne instruments which provide direct observations; however, this method is expensive and only produces a single set of data per flight. Another method to observe the  $D$  region of the ionosphere, incoherent scatter radar, is described in section 2. Data from incoherent scatter radars are limited because the technique is less sensitive than using rocket-borne instruments. In addition, incoherent scatter radars are both expensive to build and do not give a good coverage of different latitudes as there are a limited number.

Variation in electron density can also be measured using a riometer (relative ionospheric opacity meter for extraterrestrial emissions of radio noise) [Little and Leinbach, 1959]. A riometer measures the incident cosmic noise level at the Earth's surface and can be used to measure time-varying ionospheric absorption by subtracting the cosmic noise from a predefined quiet day curve [Browne *et al.*, 1995].

The absorption of radio waves is controlled by the collision frequency of thermal electrons with neutral molecules. As the radio wave travels through the ionosphere, it transfers kinetic energy to the free electrons which causes them to collide with the gas molecules in the ionosphere and thus lose their energy in the collision. If the electron does not collide with a neutral, then the kinetic radio wave energy is reradiated [Stauning, 1996]. The electron density varies spatially and temporally [Budden, 1985], causing changes in the amount of absorption that occurs, hence altering the received signal [Hunsucker and Hargreaves, 2002].

Thompson *et al.* [1986] give the maximum likely value of ionospheric absorption at 100 MHz and a zenith angle of  $60^\circ$  as 0.1 dB during the daytime and decreasing to 0.01 dB during the night. The degree of ionospheric absorption is inversely proportional to the square of the frequency; this means that absorption effects



**Figure 1.** Typical electron density ( $N_e$ ) and electron-neutral collision frequency  $\nu_{en}$  in the lower ionosphere.

are most prominent at low frequencies. The effect of absorption on radio observations can therefore be minimized by observing at higher frequencies. However, for low-frequency radio telescopes, such as LOFAR (low-frequency radio array) [van Haarlem et al., 2013] and the LWA (long wavelength array) [Ellingson et al., 2009], the absorption can affect observations that require low frequencies, for example, measurements of the sky averaged redshifted 21 cm brightness temperature [Vedantham et al., 2014; Sokolowski et al., 2015].

A basic riometer consists of a single beam fixed at zenith; however, Detrick and Rosenberg [1990] developed multibeam riometers that are capable of measuring power in multiple directions. This allows the spatial distribution of rapid variations in the level of absorption of cosmic radio noise to be observed. The riometry technique is limited by its inability to directly determine the altitudes at which the majority of absorption occurs as most riometers operate at a single frequency, typically around 30–40 MHz [Honary et al., 2011].

When using the riometry technique, the cosmic noise absorption,  $A$ , is found using

$$A(\text{dB}) = 10 \log \frac{P_0}{P}, \tag{1}$$

where  $P_0$  is the power that would be received if there was no absorption and  $P$  is the actual power received by the riometer. Sen and Wyller [1960] detail a set of formulas that relate the absorption to the radio frequency, electron-neutral collision frequency, and the gyration frequency, based on the theories of absorption devised by Lassen, Appleton, and Hartree independently in the 1920s and 1930s. This Sen-Wyller formulation indicates that the electron density  $N_e(h)$  as a function of height,  $h$ , is related to the radio wave absorption by

$$A(\text{dB}) = 4.6 \times 10^{-5} \int \frac{N_e(h)\nu_{en}(h)}{(\omega \pm \omega_g)^2 + \nu_{en}(h)^2} dh, \tag{2}$$

where  $\omega$  is the angular frequency of the radio wave,  $\nu_{en}(h)$  is the electron-neutral collision frequency, and  $\omega_g$  is the electron gyrofrequency which is defined by  $\omega_g = q_e B_0 / m_e c$ , where  $q_e$  is the charge on the electron,  $B_0$  is the magnitude of the magnetic field,  $m_e$  is the mass of the electron, and  $c$  is the speed of the light. To calculate the electron gyrofrequency, the magnetic field is approximated as a dipole, allowing a magnetic field value to be calculated.

From equation (2) it can therefore be seen that the absorption of radio waves is dependent on the electron-neutral collision frequency as a function of height as well as the electron density. While the electron density varies with incoming flux, the electron-neutral collision frequency (found using inputs generated by the NRL-MSISE-00 reference atmosphere [Picone et al., 2002]) is proportional to neutral density and therefore causes absorption to occur more at lower regions of the ionosphere, such as the  $D$  region. Typical values for  $N_e$  and  $\nu_{en}$  in the lower ionosphere are shown in Figure 1.

*Hargreaves and Friedrich* [2003] investigated whether ground-based absorption measurements at a single frequency can be used to detail the electron density in the  $D$  region, given that the relationship between radio absorption and electron density cannot be found precisely due to the spectrum of the ionizing electrons not being constant. They took data from an incoherent scatter radar, the European Incoherent Scatter Facility (EISCAT) UHF which is located in Ramfjordmoen, Norway, using Common Programme runs between 1985 and 1992 and found the median electron density values at 5 km intervals between 70 km and 100 km. This was then compared to the average absorption measurements taken from four riometers located within 10 km of the EISCAT UHF by using the relationship that the absorption at height  $h$  is proportional to the electron density, as shown by equation (2). They then used the ratio  $N_e(h)/A$  for the  $N_e(h)$  values obtained from the EISCAT data to calculate if the auroral radio absorption could be interpreted to estimate the electron density. They concluded that an estimation of approximate electron density at altitudes between 80 km and 100 km could be obtained from measurements of radio absorption; however, *Hargreaves and Friedrich* [2003] only used data from four riometers, all of which measured absorption at 27.6 MHz. If data are taken from a range of frequencies instead of just a single frequency, then the height profile of the electron density in the  $D$  region can be recovered. *Kero et al.* [2014] used the Kilpisjärvi Atmospheric Imaging Receiver Array (KAIRA) as a multifrequency riometer, measuring absorption at a range of frequencies from 10 to 80 MHz. They found that multifrequency riometry measurements can provide realistic electron density height profiles when using the Delayed Rejection Adaptive Metropolis-Hastings Markov chain Monte Carlo (MCMC) algorithm and a model of the electron precipitation that occurs. Electron precipitation is caused by very low frequency waves colliding with electrons in the radiation belt and causing them to propagate into the ionosphere, thus changing its shape and conductivity.

This paper explores a new method of using multifrequency absorption data to infer an electron density height profile for the ionosphere using IONONEST. IONONEST uses the nested sampling method implemented through the publicly available MULTINEST software library [*Feroz and Hobson, 2008; Feroz et al., 2009*]. It is a tool for finding the shape of the electron density height profile from multifrequency measurements of the absorption, and it provides a general framework for full Bayesian evaluation of different electron density models, including quantitative model comparison. This enables fast recovery of model parameters which can be used for time-dependent studies of the lower ionosphere at a range of latitudes, allowing models of the currently under observed  $D$  region to be improved.

In this paper we investigate the validity of two different electron density height profile models by fitting them to electron density height profiles obtained by an incoherent scatter radar, the EISCAT VHF radar, before using IONONEST to fit electron density height profiles to measured absorption data from KAIRA. Section 2 details the instruments used and observations taken, while section 3 introduces nested sampling, as well as the models used to fit the electron density height profiles. Section 3 also investigates whether realistic electron density height profiles such as those observed by the EISCAT VHF can be recreated using these models. Section 4 describes the implementation of IONONEST and how it uses the nested sampling method to recover electron density height profiles from absorption measurements, as well as detailing the application of IONONEST to simulated data sets to establish if IONONEST can return true electron density height profiles from absorption data. Section 5 details the application of IONONEST to KAIRA absorption data and compares the resulting electron density height profiles to those obtained using the EISCAT VHF.

## 2. Observations

In March 2015, 6 h of observations were taken using both KAIRA and the EISCAT VHF. These two instruments are located approximately 85 km apart. Here we detail the instruments and the observations made with them.

### 2.1. KAIRA

The Kilpisjärvi Atmospheric Imaging Receiver Array (KAIRA) is located in Kilpisjärvi, Finland (longitude = 20.76°E, latitude = 69.07°N). It is a radio telescope that uses the antenna system and signal processing of a LOFAR telescope [*van Haarlem et al., 2013*] to form multiple beams on the sky [*Vierinen et al., 2013; McKay-Bukowski et al., 2015*].

KAIRA consists of a high-band antenna (HBA) array and a low-band antenna (LBA) array, each made up of 48 dual polarization antennas and can be used as a multibeam, multifrequency riometer [*Kero et al., 2014; McKay-Bukowski et al., 2015*]. The antennas in the LBA array are crossed inverted V dipole aeriels and are

scattered quasi-randomly in a circle of diameter 34 m. The HBA array is made up of 48 tiles which each contain 16 crossed bowtie antenna elements. The KAIRA HBA array operates at frequencies from 110 MHz to 270 MHz, while the LBA array operates at frequencies from 10 MHz to 80 MHz. The typical sampling clock frequency used is 200 MHz. This affects the width of the subbands which are created by a polyphase filter bank splitting the received signal into 512 subbands. When using 16 bit sampling the HBA and LBA arrays can select up to 244 of these subbands, allowing up to 244 different pointing directions coupled with a frequency. These are known as beamlets. For 4 or 8 bit sampling, 976 or 488 beamlets can be used, though by reducing the sampling, the dynamic range worsens.

The power that KAIRA measures is dependent on the total system temperature,  $T_{\text{sys}}$ .  $T_{\text{sys}}$  is the noise from the whole system and includes both the noise measured from the sky ( $T_{\text{sky}}$ ) and noise due to the receiver system,  $T_{\text{instrument}}$ .

$$T_{\text{sys}} = T_{\text{sky}} + T_{\text{instrument}}. \quad (3)$$

The noise due to the receiver system is caused by thermal noise in the electronic components that make up the array, caused by thermal agitation of free electrons and typically has a noise spectrum that is flat with respect to frequency. This is known as “white noise.” The noise in each component is normally independent of noise from the other components, and the power due to the noise is given by

$$P = k_B T. \quad (4)$$

The total noise of all the components is estimated by adding together all the individual contributions. This is the instrument noise,  $T_{\text{instrument}}$ , and is calculated using the system equivalent flux density which has been measured for KAIRA by *McKay-Bukowski et al.* [2015]. This allows us to calculate the true power received from the sky,  $P_{\text{sky}}$ , as we can subtract the instrumental noise from the total power received by KAIRA,  $P_{\text{sys}}$  using

$$P_{\text{sys}} = B (kS + P_{\text{sky}}), \quad (5)$$

where  $B$  is the bandpass,  $S$  is the system equivalent flux density, and  $k$  is a conversion factor which is needed to convert the system equivalent flux density from kJy to Analog Digital Units (ADU), assuming the gain is constant, because  $P_{\text{sys}}$  and  $P_{\text{sky}}$  are both measured in ADU.

KAIRA suffers from both broadband and narrowband radio frequency interference (RFI). A moving window filter is passed over the data in order to remove RFI. The RFI contaminated data are flagged, allowing for the removal of entire time periods or entire frequency bands if the time period or frequency is too heavily contaminated. Otherwise, the data are filtered out using the average of the moving window filter. Below 20 MHz, RFI caused by short-wave radio communications dominates [*McKay-Bukowski et al.*, 2015].

KAIRA also suffers from a certain type of RFI known as “Dragon’s teeth” (see Figure 2). These manifest in the data as broadband RFI which is triangular in shape. The source of this RFI has yet to be determined.


The LBA array at KAIRA was used for absorption measurements in order to exploit the multifrequency dependence of the absorption signal. The observations began at 23:00 UT on 1 March 2015 and ended at 05:00 UT on 2 March 2015. These observations were made in the early morning sector as local sunset at KAIRA was at 15:39 UT and local sunrise was at 05:58 UT.

Eight pointing directions were used with the KAIRA LBA array, with half the beamlets pointing at zenith and observing between 9.77 MHz and 80.66 MHz, with a frequency spacing of 0.5859 MHz. Twenty of the remaining beamlets were pointing at the north celestial pole and observing between 9.77 MHz and 76.56 MHz with frequency spacings of 3.5156 MHz, and the other 102 beamlets had static pointings above the EISCAT site, equally split between six pointing directions and observing between 9.77 MHz and 75.39 MHz, with frequency spacings of 4.1016 MHz. Table 1 lists the elevation angle of these pointing directions and corresponding altitude over the EISCAT VHF. The range of the full-width half maxima at zenith of the KAIRA LBA beams used ranged between 8° at 80 MHz to 40° at 10 MHz. The average power for each beamlet and each polarization were recorded with the KAIRA LBA array at a temporal resolution of 1 s.

Below we detail the methods involved in using KAIRA as a riometer to obtain absorption data.

### 2.1.1. Quiet Day Surface

Standard practice in riometry is to define a quiet day curve (QDC) which represents the power level that is received on a day when minimal absorption occurs and the ionosphere is assumed to be transparent.



**Figure 2.** (top) The quiet day surface created for KAIRA (beam pointing at zenith) using observations from 1 to 8 March 2015. (middle) The observed power at KAIRA for a beam pointing at zenith and observing between frequencies 9.77 MHz and 80.66 MHz on 1 to 2 March 2015. (bottom) The absorption,  $A$ , as measured by KAIRA for a beam pointing at zenith. The absorption is found using equation (1) with the values in Figure 2 (top) as  $P_0$  and the values in Figure 2 (middle) as  $P$ . The vertical red lines in this absorption plot show peaks in absorption. RFI can be seen to dominate at all frequencies below 16.80 MHz and “Dragons teeth” RFI (as detailed in section 2.1) can be clearly seen in Figures 2 (middle) and 2 (bottom) at 01:00 UT and 03:30 UT, respectively.

**Table 1.** KAIRA Pointing Directions and Corresponding Altitude Above the EISCAT VHF<sup>a</sup>

Pointing Direction	Elevation (deg)	Azimuth (deg)	RA	Declination	Altitude Above EISCAT VHF (km)	Beamlets	Frequency Range (MHz)
1	90 (zenith)	180	-	-	-	122	9.77–80.66
2	-	-	180	90	-	20	9.77–76.56
3	35	313.95	-	-	59.5	17	9.77–75.39
4	40	313.95	-	-	71.3	17	9.77–75.39
5	45	313.95	-	-	85.0	17	9.77–75.39
6	50	313.95	-	-	101.3	17	9.77–75.39
7	55	313.95	-	-	121.4	17	9.77–75.39
8	60	313.95	-	-	147.2	17	9.77–75.39

<sup>a</sup>Multiple frequencies are observed using KAIRA between the ranges displayed.

By comparing the signal received by the riometer with a QDC, the extent to which the background noise is being absorbed by the ionosphere can be found [Browne *et al.*, 1995]. Therefore, the difference between received signal power and the corresponding QDC value directly represents absorption. The QDC is needed because the rotation of the Earth and the apparent motion of the radio sky cause variation in the power received by the riometer. The Sun also causes increases in the received power as well as ionization in the *D* region at a height of about 80 km during daylight hours. This periodic variation of absorption in the ionosphere is dependent on the Sun's zenith angle and thus also contributes to the characteristic shape of the QDC. Such curves need generating regularly throughout the year, as well as throughout the 11 year solar cycle.

Here we follow the method described in Browne *et al.* [1995] for generating a QDC, developed for IRIS (Imaging Riometer for Ionospheric Studies). Data are arranged into sidereal days, and 10 min averages are calculated over each day. Each time interval is then sorted in descending order, and the mean of the second and third highest values for each time interval is calculated. This mean is assumed to be a good estimate of quiet time absorption for that time interval in the QDC. Once a value for each time interval has been calculated, interpolation can be performed to obtain quiet values for each second of the sidereal day.

Unlike a standard riometer, we used KAIRA to simultaneously observe with 244 beamlets, so that instead of a quiet day curve, a quiet day *surface* is created. To create the quiet day surface, the variation in power level is measured throughout the day at different frequencies for each beam pointing in Table 1. An example of a quiet day surface for the KAIRA beam pointing at zenith is shown in Figure 2 (top). For the creation of this quiet day surface observations were made between 1 March 2015 and 8 March 2015. The quiet day surface provides a value for  $P_0$ , as required to calculate the absorption using equation (1).

### 2.1.2. Absorption Surface

Using equation (1) (with  $P_0$  given by the frequency-dependent power in the quiet day surface shown in Figure 2 (top) for the zenith case), the absorption value for each time, frequency and pointing direction can be calculated from the measured power on the day we are observing,  $P$ . Figure 2 (middle) shows the measured power,  $P$ , for the zenith pointing between 23:00 UT on 1 March 2015 and 05:00 UT on 2 March 2015. The resulting absorption surface is shown in Figure 2 (bottom).

The absorption data calculated with equation (1) are dominated by RFI below 16.80 MHz; these subbands are therefore removed from the data set. The remaining data are smoothed using the Hanning smoothing algorithm to mitigate any further effects from RFI. This uses a triangular smoothing kernel to calculate a running mean along the frequency axis, with the central subband weighted by 0.5 and the two adjacent subbands weighted by 0.25.

## 2.2. EISCAT

EISCAT was built in 1981 and is a network of incoherent scatter radars which are used to investigate the Earth's atmosphere and ionosphere [Folkestad *et al.*, 1983]. Incoherent scatter causes ionospheric electrons to start to oscillate due to the electric field of the wave transmitted by the radar. This oscillation causes them to reradiate, giving rise to a weak, but detectable, signal whose frequency spectrum can be directly related to the key parameters of the scattering plasma. In this paper we use the monostatic EISCAT VHF radar which is located at Ramfjordmoen (near Tromsø, Norway, longitude = 19.23°E, latitude = 69.38°N). The EISCAT VHF radar operates in the 224 MHz band [Baron, 1986] and has a beam width of  $0.6^\circ \times 1.7^\circ$ .

The EISCAT VHF obtains height profiles of the electron density in the ionosphere by transmitting coded pulses of electromagnetic energy. These pulses interact with electrons in the ionosphere via Thomson scattering [Aksnes *et al.*, 2006]. The energy is reradiated from these electrons; however, as the electrons all have individual thermal motion the received signal will be incoherent; therefore, the received signal must be decoded in order to determine the power spectrum. Using the shape of this spectrum, the electron density, electron temperature, ion temperature, and line-of-sight ion drift can be found [Rishbeth and van Eyken, 1993].

The EISCAT VHF radar has a range of observing modes that enables it to be optimized for the relevant geophysical process being studied. It is capable of observing in the altitude range of 50–2500 km and can have a temporal resolution of less than one second. The limiting value of electron density for the EISCAT VHF occurs when the level of random error in an estimate of a plasma parameter becomes too high and the threshold of detectability is reached because the measurement becomes meaningless. This generally occurs at altitudes below 80 km. At higher altitudes (>120 km), the resolution in altitude worsens to over 1 km when using the radar-observing mode “manda” (see <https://www.eiscat.se/about/experiments2/experiments> for details).

The EISCAT VHF observations were made with a temporal resolution of 1 min and a pointing direction of zenith. The radar-observing mode manda was used for the observations as it is optimized for the measurement of plasma parameters in the *D* and *E* regions of the ionosphere, the heights at which absorption primarily occurs. The observations began at 23:00 UT on 1 March 2015 and ended at 05:00 UT on 2 March 2015 (local sunset at the EISCAT VHF was at 16:44 UT and local sunrise was at 06:13 UT).

The observations made using the EISCAT VHF are used throughout this paper for verification that the models of electron density used are capable of returning realistic electron density height profiles, as well as for comparison to information derived from the KAIRA data, allowing us to validate our methods.

### 3. Height Profile Determination

To create electron density height profiles from the absorption data, a model is required for the electron density height profile. In this section we introduce two electron density height profile models. By varying the parameters in the models, we can simulate profiles which can then be compared to observed data. To do this, a method called nested sampling [Skilling, 2004] is used. This method is based on Bayesian inference, which allows us to estimate a set of parameters,  $\theta$ , in a model,  $M$ , given data,  $x$ . It uses Bayes theorem, which states,

$$p(\theta|x, M) = \frac{p(x|\theta, M)p(\theta|M)}{p(x|M)}, \quad (6)$$

where the posterior probability distribution is  $p(\theta|x, M) \equiv P(\theta)$ , the likelihood is  $p(x|\theta, M) \equiv \mathcal{L}(\theta)$ , and the prior probability distribution is given by  $p(\theta|M) \equiv \Pi(\theta)$ .  $p(x|M)$  is the Bayesian evidence, which is the factor that normalizes the posterior over the prior volume and can be ignored in the case of parameter estimation. In the case of model selection the evidence is needed in order to determine which model performs best when using a common data set.

#### 3.1. Nested Sampling

While most model fitting algorithms concentrate on parameter estimation, calculating the Bayesian evidence often as a by-product, the nested sampling technique calculates the Bayesian evidence directly. The nested sampling algorithm selects a range of points from the parameter space to create models; these models are then compared to the data, and a likelihood value is returned. A weight is then assigned to the point in the parameter space with the lowest likelihood before it is recorded and then discarded from the selection of points in the parameter space, and a new point with a higher likelihood is selected from the parameter space using a Monte Carlo method. This causes the sampler to converge on the likelihood peaks of the parameter distribution. Our knowledge of the model parameters is defined by the priors before the model is fitted to the data. A prior can be defined using the posterior of previous observations or as the available parameter space.

MULTINEST is a publicly available algorithm developed by Feroz, Hobson and Bridges [Feroz and Hobson, 2008; Feroz et al., 2009], implementing the nested sampling method. It uses Bayesian inference [e.g., Trotta, 2008] to select a model using the evidence as the factor to normalize the posterior. This is done by transforming the multidimensional evidence integral into a one-dimensional integral in order to calculate the evidence.

This method reduces the computational time required to fully explore the parameter space and is effective for posteriors which are multimodal. MULTINEST has recently been improved to also use importance nested sampling [Cameron and Pettitt, 2013] which uses all sampled points in the evidence calculation. To date, MULTINEST has been used for a number of astronomical applications, including dispersion measure fitting (TEMPONEST [Lentati et al., 2014]) and object detection [Feroz et al., 2008].

#### 3.2. *D* Region Models

In this section we detail the *D* region electron density height profile models that have been chosen to be investigated in this paper. The two models chosen are Model 1, a two parameter exponential model, and Model 2, a polynomial model.

These models and the methods used to evaluate the strength of evidence for them are discussed below. For each model, the priors are stated.

##### 3.2.1. Model 1: Two Parameter Exponential Model

Wait and Spies [1964] detailed a simple two parameter exponential model which was determined to be the best exponential fit for broadband measurements using VLF radio emissions from lightning to profiles



obtained by rocket measurements by *Cheng et al.* [2006]. The model is of the form

$$N_e(h) = 1.43 \times 10^7 \exp[-0.15h'] \exp[(\beta - 0.15)(h - h')], \quad (7)$$

where  $h'$  controls the height of the profile and  $\beta$  controls the sharpness of the ionospheric transition (the change of  $N_e$  with altitude).  $\beta$  can vary between 0 and 1, while we allow  $h'$  to vary from 60 km to 110 km. *Thomson* [1993] also used this model to compare VLF theory to measurements and discuss values over which the two parameters,  $h'$  and  $\beta$ , can vary for daytime observations.

### 3.2.2. Model 2: Polynomial Model

*Parthasarathy et al.* [1963], *Stoker* [1987], and more recently, *McKay-Bukowski et al.* [2015] have all fitted a polynomial model to the electron density in the  $D$  region. Here we also investigate the possibility that the  $D$  region can be fitted by a polynomial of the form

$$N_e(h) = \exp\left(a_0 + a_1 \left(\frac{h - h_{\text{ref}}}{50}\right) + \dots + a_n \left(\frac{h - h_{\text{ref}}}{50}\right)^n\right). \quad (8)$$

The parameters ( $\theta = a_0, a_1, \dots, a_n$ ) are allowed to vary over a wide range of standard uniform priors, where  $-200 \leq \theta \leq 200$ . The reference height,  $h_{\text{ref}}$ , is set equal to 80 km, which is the height at which  $N_e(80) = \exp(a_0)$ .

### 3.2.3. Comparison of Models

The evidence ratios (Bayes factors) of the model fits can be used to determine the optimal fit using the Jeffreys scale [*Jeffreys*, 1961]. The Jeffreys scale provides a formal way of evaluating which model is better supported by the data.

Model 2 can be compared to Model 1 using the Bayes factor

$$B_{12} = \frac{p(x|M_1)}{p(x|M_2)} = \frac{\int p(\theta_1|M_1) p(x|\theta_1, M_1) d\theta_1}{\int p(\theta_2|M_2) p(x|\theta_2, M_2) d\theta_2}, \quad (9)$$

where  $x$  is the observed data and the two different models are  $M_1$  and  $M_2$  with parameters  $\theta_1$  and  $\theta_2$ . This Bayes factor gives the change in relative probability of the two models. If the Bayes factor,  $B_{12}$ , is greater than 1, then Model 1 is favored over Model 2 and vice versa if  $B_{12}$  is less than 1. We use a revised version of the Jeffreys scale as detailed in *Trotta* [2008].

The quality of the prediction can also be evaluated using the standardized mean squared error [*Rasmussen and Williams*, 2005]. The standardized mean square error (SMSE) is calculated by computing the squared residual between the mean prediction and the target at each point; the mean of these values is then found (the mean squared error), and the value is then normalized by the variance of the data to output the SMSE,

$$\text{SMSE} = \frac{\sum_{i=1}^n (Z_i^* - Z_i)^2}{Z_{\text{var}}}, \quad (10)$$

where  $Z_i^*$  is the estimated parameter,  $Z_i$  is the measured parameter, and  $Z_{\text{var}}$  is the variance of the measured parameter. Smaller values obtained by the SMSE indicate a better goodness of fit between the estimated and measured parameters.

## 3.3. Model Fits to EISCAT Data

To validate the models described above, we fit them directly to the EISCAT height profiles using the nested sampling method. This enables a comparison of the electron density height profiles output by the models to the EISCAT observations to be made, enabling us to check whether the models are capable of producing realistic electron density height profiles. The strength of evidence is then assessed for each of the models using a modified Jeffreys scale [*Trotta*, 2008] as well as the SMSE.

We begin by establishing the order of the polynomial needed to fit the data for Model 2. To do this, we vary the order of the polynomial that we use to establish at what point the evidence reaches its maxima. This process is shown in Table 2. We find that a second-order polynomial creates the best fitting model fitted to a typical electron density height profile. While the calculation of the SMSE (equation (10)) will always show a better fit with increase in order of the polynomial, the Bayesian fitting integrates over all parameters in each model and so includes a penalty for complex models that include additional unnecessary parameters (Occam's razor). A complex model should only return a larger posterior probability if it needs to be that complex to fit the data.

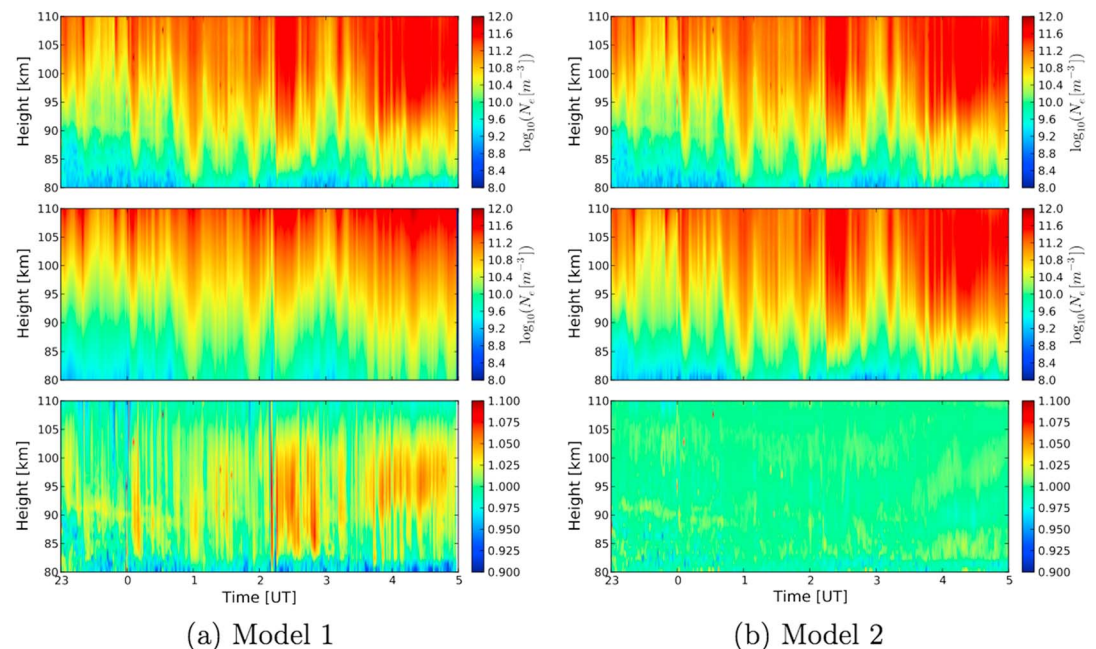
**Table 2.** The Order of the Polynomial Fit to a Typical EISCAT Height Profile With the Fitted Maximum Likelihood Polynomial Coefficients, as Well as the Natural Logarithm of the Evidence for a  $3\sigma$  Prior Volume ( $\ln Z$ ) and the Difference Between These ( $\Delta \ln Z$ )

Order	$a_0$	$a_1$	$a_2$	$a_3$	$\ln Z$	$\Delta \ln Z$	SMSE
1	$9.66 \pm 0.04$	$5.81 \pm 0.09$	–	–	$-271.28 \pm 0.07$	–	20.73
2	$8.90 \pm 0.08$	$16.71 \pm 0.51$	$17.19 \pm 0.73$	–	$-13.04 \pm 0.06$	$-258.24$	0.30
3	$8.97 \pm 0.10$	$15.36 \pm 1.24$	$-12.07 \pm 4.33$	$-5.26 \pm 4.36$	$-31.19 \pm 0.15$	18.15	0.27


Figure 3 shows the model fits to the height profiles obtained using EISCAT, while Figure 4 shows comparisons of the height of the peak electron density and the peak electron density output by the models to that measured by EISCAT between 80 and 110 km. We find that Model 1 returns increases in electron density that correspond to increases in the EISCAT data. The model does not, however, return the peaks in electron density at the same height as in the EISCAT data nor does it return the true peak electron density when the ionosphere is particularly active. As Model 1 is an exponential model the peak electron density will always be at the highest point in the altitude range, thus giving a constant peak electron density of 110 km in Figure 4a. The peak electron density is always overestimated because in attempting to fit to the rest of the electron density profile Model 1 always overshoots at the top of the altitude range. In fact, *Cheng et al.* [2006] only apply this model to data below 90 km, possibly for this reason. Model 2, however, succeeds to some degree in returning the height of the peak electron density in the EISCAT data, as well as returning the peak electron density, as shown in Figure 4b.

Figure 3a (bottom) is the ratio between the EISCAT electron density height profiles and Model 1 fitted to the EISCAT electron density height profiles. It shows that as well as being overestimated at higher altitudes in the height range, this model causes the electron density to also be consistently overestimated at the lower altitudes (80–82 km) but underestimated in the middle of the height range when the ionosphere is particularly active, a problem that is not seen in the Model 2 fit as Model 2 is flexible enough to accommodate almost any profile shape. Model 2 thus returns a better representation of the EISCAT data.

We evaluate the fits using the methods described in section 3.2.3. We calculate the Bayes factor, finding that  $B_{12} = 0.196$ ; therefore, Model 2 is favored over Model 1. The value of  $|\ln B_{21}|$  is found to be 2.25. Using the



**Figure 3.** (a) Model 1 fitted to EISCAT data. (b) Model 2 fitted to EISCAT data. (top row) EISCAT height profile. (middle row) Model fitted to EISCAT data. (bottom row) The ratio of the EISCAT height profiles and the model fitted to the EISCAT height profiles. The color bars are saturated at the extremes in order to emphasize the features.



**Figure 4.** Comparison of the height of the peak electron density and the peak electron density output by (a) Model 1 and (b) Model 2 to that measured by EISCAT between 80 and 110 km.

empirical scale given in Table 1 of Trotta [2008], this means that there is “weak evidence” for Model 2 being favored over Model 1.

To compare the Model 1 and Model 2 fits to the EISCAT electron density height profiles using the SMSE,  $Z_i$  is set to be the EISCAT electron density,  $Z_{\text{var}}$  to be the variance in the EISCAT data, and  $Z_i^*$  to be the Model 1 or Model 2 fit to the EISCAT electron density. For Model 1 the SMSE value was found to be 0.76, while for Model 2 it was found to be 0.11. This shows that Model 2 is a better fit to the EISCAT electron density height profiles as the SMSE value obtained is smaller than that found for Model 1, indicating a better goodness of fit between the data estimated by Model 2 and the EISCAT data. This is in agreement with the Bayes factor calculated above.

#### 4. IONONEST

In order to recover electron density height profiles from the absorption data observed by KAIRA, we implement the nested sampling method detailed in section 3.1 by interfacing with MULTINEST. We have named this implementation IONONEST.

IONONEST is a tool that can find the shape of the electron density height profile from multifrequency measurements of absorption. It assumes a model for the electron density height profile and varies the parameters in the model. The absorption profile for this modeled electron density height profile is then created using equation (2) and compared to the absorption data observed by KAIRA.

The inputs to IONONEST are an observed, or simulated, data set of absorption values, an electron density height profile model and priors, where the priors are the parameter space over which parameters in the electron density height profile models are allowed to vary. MULTINEST then selects values for the parameters from the priors and IONONEST creates an electron density height profile and calculates the associated absorption values at the observing frequencies. It then compares the calculated absorption values to the observed or simulated absorption values and returns a likelihood to MULTINEST. At any point, MULTINEST has 1000 live points, and once IONONEST returns the likelihood it has calculated to MULTINEST, MULTINEST drops the point with the lowest likelihood from the set of live points and selects new parameter values in order to replace it. This process terminates when MULTINEST finds the log evidence to a precision of 0.1.

The noise on the KAIRA data is expected to be Gaussian, and we assume uncorrelated noise between frequency channels. The log-likelihood that the data are described by the electron density height profile model is therefore given by

$$\mathcal{L} = \sum_{i=1}^n \left( -\frac{1}{2} \ln \left( 2\pi\sigma_{\text{obs},i}^2 \right) - \frac{(A_{\text{obs},i} - A_{\text{mod},i})^2}{2\sigma_{\text{obs},i}^2} \right), \quad (11)$$

where  $A_{\text{mod}}$  is the absorption due to the electron density height profile model,  $A_{\text{obs}}$  is the observed absorption, and  $\sigma_{\text{obs}}$  is the uncertainty associated with the observed absorption values. For parameter estimation, we can ignore the first term in this sum.

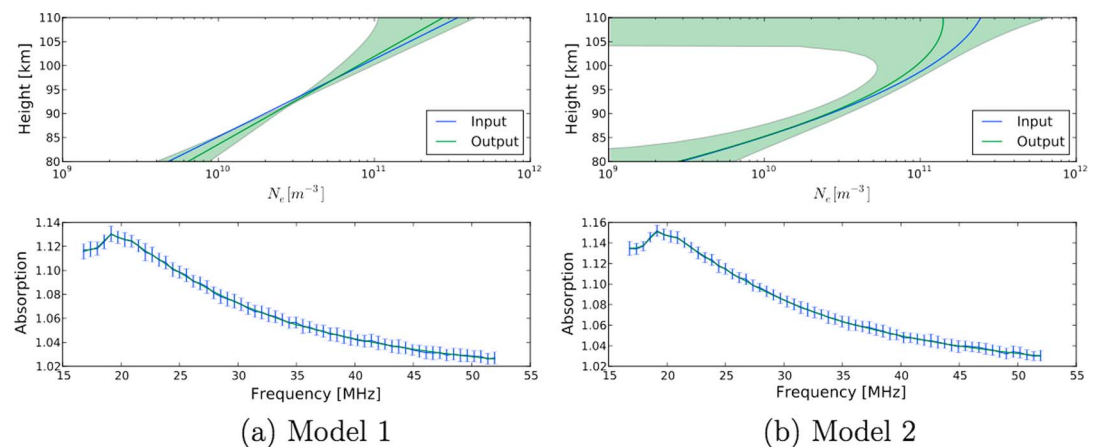
**Table 3.** Values for the Input Parameters That Created the Simulated Data and the Parameters Output for the Best Fit to the Simulated Data

Parameter	Input	Output
<i>Model 1</i>		
$h'$	66.351	$64.474 \pm 3.453$
$\beta$	0.293	$0.276 \pm 0.034$
$\log(k)$	-5.088	$-5.083 \pm 0.027$
<i>Model 2</i>		
$a_0$	7.977	$7.934 \pm 0.964$
$a_1$	12.858	$13.359 \pm 5.513$
$a_2$	-9.127	$-11.389 \pm 9.418$
$\log(k)$	-5.088	$-5.087 \pm 0.030$

Before applying this method to KAIRA absorption data, we first verify that it is possible to return accurate parameter estimates by simulating absorption data and using IONONEST to find fits to that data.

#### 4.1. Fit to Simulated Absorption Data

Having shown in section 3.3 that both models provide reasonable fits to electron density height profiles derived from EISCAT data, we now run our IONONEST code on absorption data sets simulated from generated height profiles. From this test it can be seen if the parameters input into the simulation can be returned by IONONEST. We simulate a single height profile using each of the models detailed in section 3.2. These electron density profiles are simulated using the input parameters detailed in Table 3 and are shown in Figure 5 (top) (labeled with "input"). Absorption values for frequencies in the range 17.38–56.05 MHz are calculated for the height profiles, and white noise is added as well as the instrumental, frequency-dependent noise described in section 2.1 at a similar level to that of a KAIRA observation. This is done by adding another parameter to fit with IONONEST,  $k$ , as used in equation (5). The simulated absorption data corresponding to the simulated electron density profiles are shown in Figures 5a (bottom) and 5b (bottom). These simulated absorption profiles are then input into IONONEST. The priors are set to vary between the maximum and minimum parameter values obtained for each model in the fits to the EISCAT data shown in Figure 3  $\pm 3\sigma$ . The prior for  $k$  for Models 1 and 2 is set to a standard uniform prior,  $-5.5 \leq \log(k) \leq -4.5$ . These priors are listed in Table 4. IONONEST returns the parameters to create the height profile once it fits an absorption profile to the simulated absorption data, as shown in Figures 5a (bottom) and 5b (bottom). Figure 5 shows that IONONEST fits the simulated absorption data using Model 1 and Model 2 to within  $1\sigma$ .



**Figure 5.** (top) Simulated (a) Model 1 and (b) Model 2 height profile (“input”), with a height profile fitted to absorption data generated from this simulated height profile plotted over the top (“output”). The area enclosed by the shaded green region indicates the electron density allowed at each height by  $1\sigma$  uncertainties on the parameters given in Table 3. (bottom) Simulated absorption data with error bars ( $\pm 1\sigma$ ) plotted, with the fit to the absorption data using (a) Model 1 and (b) Model 2 as found by IONONEST plotted over the top.

**Table 4.** Priors for the Two Models Found Using  $\mu \pm 3\sigma$  as Obtained From Fitting the Models Directly to the EISCAT Data

Parameter	Minimum	Maximum
<i>Model 1</i>		
$h'$	55.65	74.71
$\beta$	0.22	0.34
$\log(k)$	-5.5	-4.5
<i>Model 2</i>		
$a_0$	5.97	10.17
$a_1$	2.25	24.81
$a_2$	-28.36	5.42
$\log(k)$	-5.5	-4.5

The parameters output by IONONEST and used to create these profiles are listed in Table 3 and labeled as “output” in Figure 5. It can be seen from Table 3 that IONONEST returns the parameter values to within  $1\sigma$  for both Model 1 and Model 2. We note that the posterior probability distributions for these parameters are highly degenerate causing the large errors seen in Figures 5a (top) and 5b (top). This results in different electron density height profiles being reproduced by IONONEST regardless of the close fit in absorption data. Model 2 has trouble reproducing the input profile above an altitude of 100 km as the uncertainties become large at altitudes greater than 100 km. The electron density height profiles returned are, however, similar to the simulated electron density height profiles, demonstrating that height profiles can be recovered from simulated absorption data.

## 5. Application to KAIRA Data


In section 4 we showed that both models are capable of recreating electron density height profiles from simulated absorption data. We now apply IONONEST to the absorption profiles obtained from the pointings at zenith and the pointings intersecting with the EISCAT VHF beam at points A and B in Figure 6. The electron density height profiles output by IONONEST are then compared to the profiles as measured by EISCAT.

### 5.1. Pointings Above KAIRA (Zenith)

First, the absorption profiles obtained from the zenith pointing are input into IONONEST. Figure 7 shows the resulting height profiles obtained by IONONEST using Model 1 and Model 2. Figures 7a (bottom) and 7b (bottom) shows the ratio between these height profiles and the EISCAT VHF height profile. Both Model 1 and Model 2 show some agreement with the EISCAT VHF electron density profiles, generally recreating the shape of the electron density variation with time. Similarly to Figure 3a, the ratio in Figure 7a (bottom) shows overestimation by Model 1 of the electron density at the top and bottom of the height range. It also shows an underestimation of the electron density by Model 1 in the middle of the height range. Figure 8a shows that the peak electron density is reasonably well recreated by Model 1; however, due to the nature of the model, the height of the peak electron density is always at the maximum of the height range being modeled. On the other hand, when lower heights of the peak electron density are seen in the EISCAT data, lower heights of the peak electron density are also seen in the electron density height profiles created by Model 2. However the peak electron density for Model 2 (as seen in Figure 8b) is smaller than that observed by the EISCAT VHF.

The electron density height profiles returned by Model 1 and Model 2 show an underestimation of electron density at the beginning and end of the observation, as well as an overestimation in electron density shown midway through the observation. Earlier in this paper it was shown that IONONEST is capable of returning simulated height profiles using these models, and so we consider that these are real features that differ from the EISCAT data because of physical differences between the two instruments.

These physical differences include the positions and differences in beam size of the two instruments. EISCAT and KAIRA are located 85 km apart and therefore are not observing the same volume of ionosphere. The EISCAT VHF beam is also very narrow, with a beam size of  $0.6^\circ \times 1.7^\circ$  whereas KAIRA has a beam size of  $10^\circ - 20^\circ$  (depending on frequency). This means that events that affect the electron density of the D region such as thin filaments due to electron precipitation may fill the EISCAT beam and appear as bright absorption events



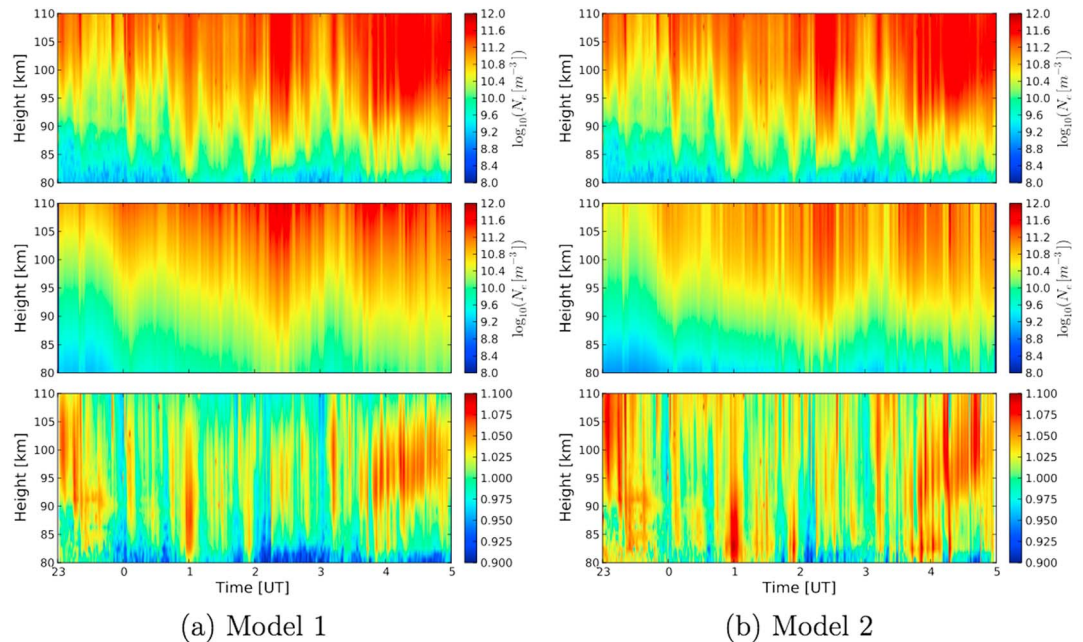
**Figure 6.** The EISCAT VHF and KAIRA pointings. The KAIRA pointings are detailed in Table 1. The lowest intersection with the EISCAT beam is at 59.5 km and the highest is at 147.2 km. Only two of these pointings (elevations of 45° and 50°, labeled as A and B) intersect with the EISCAT VHF beam within the 80–110 km height range above the Earth’s surface. We concentrate on these two pointings as the absorption is highest at these heights. We cannot fit below 80 km because the data from the NRL-MSISE-00 reference atmosphere are inaccurate below this height. Access to an ion chemistry model would enable us to extend down to lower heights.

causing a peak in the electron density. In the KAIRA observations, however, the bright absorption event caused by the thin filament will be averaged out over the large beam, and therefore, only the averaged electron density over the whole beam is output in the electron density height profiles. We do not expect to be able to capture any compact features, such as bright absorption features caused by thin filaments, to the same effectiveness as the EISCAT VHF due to the averaging effect over the KAIRA beam.

Figure 8 shows that the electron density height profiles obtained by Model 1 show a similar peak electron density to the EISCAT data, while Model 2 shows consistently lower peak electron densities. However, in Figure 4a, we demonstrated that a Model 1 fit to the EISCAT data overestimated the peak electron density. Bearing in mind the smaller size of the EISCAT beam, we might assume that if Model 1 was a more realistic model, then it should also underestimate the data.

Using the quality of fit assessment outlined in section 3.2.3, we can evaluate the quality of the prediction using the SMSE (equation (10)). When calculating the SMSE,  $Z_i^*$  is the estimated parameter, which is the absorption fit to the KAIRA data, while  $Z_i$  is the measured parameter, which is the absorption data observed by KAIRA.  $Z_{var}$  is the variance of the measured parameter, and so in this case it is the variance of the absorption data observed by KAIRA.

The SMSE values obtained for Model 1 and Model 2 are 0.21 and 0.35, respectively. While the lower value for the Model 1 estimated absorption fit to the data implies that it is an improved fit when compared to Model 2, reinforcing what can be seen in the ratio plots in Figure 7, both models are reasonable fits to the data.



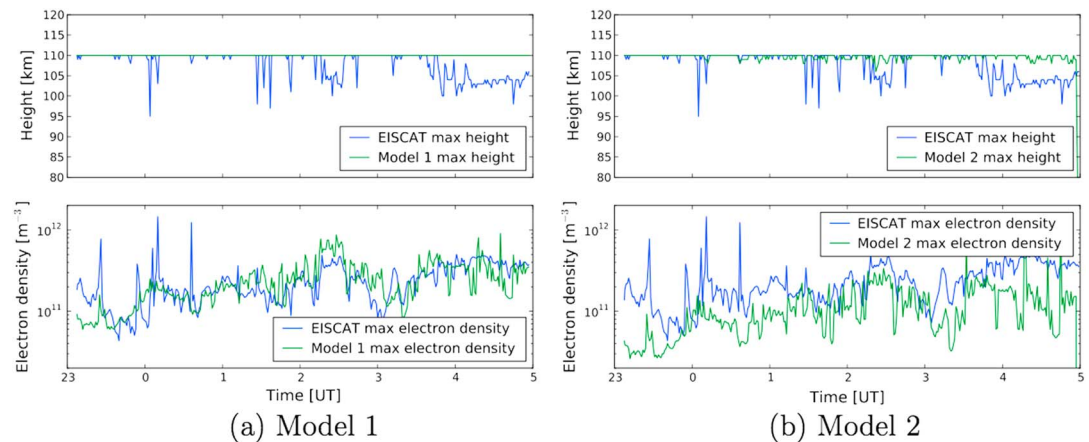
**Figure 7.** These plots show a comparison between the height profiles generated using (a) Model 1 and (b) Model 2 fitted to the KAIRA absorption data for the zenith pointing and the height profiles obtained from EISCAT. (top row) The EISCAT height profiles for 23:00 UT on 1 March 2015 until 05:00 UT on 2 March 2015. (middle row) The model height profiles fitted to the KAIRA absorption data for the same time period and pointing. (bottom row) The ratio between the EISCAT data and the fit to KAIRA absorption data. The color bars are saturated at the extremes in order to emphasize the features.

We also evaluate the fits by calculating the Bayes factor, finding that  $B_{12} = 1.087$ . This again shows that Model 1 is favored over Model 2; however, the value of  $|\ln B_{12}|$  is found to be 0.10. According to the empirical scale given in Table 1 of Trotta [2008], this means that the evidence is “inconclusive” for Model 1 being favored over Model 2.

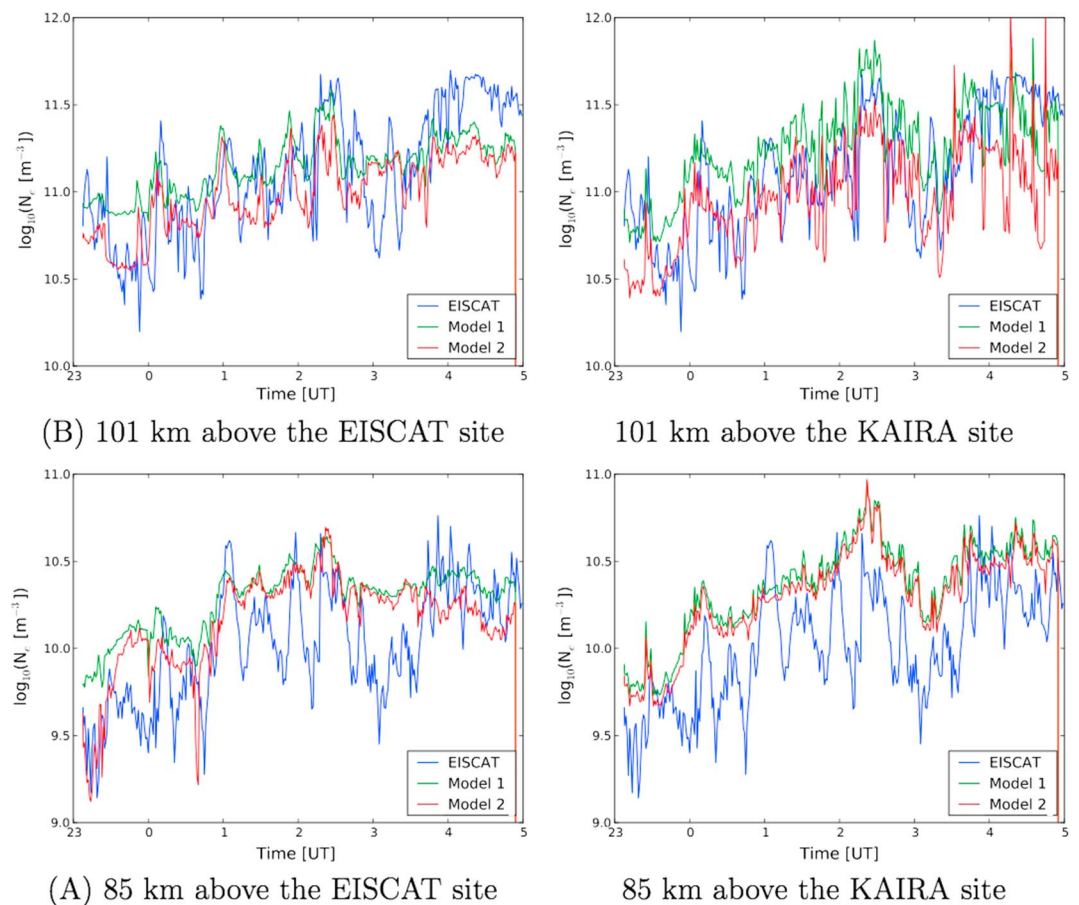
**5.2. Pointings Above EISCAT**

We now use the KAIRA absorption data for pointings A and B in Figure 6, enabling the two instruments to observe the same volume of ionosphere, albeit with KAIRA’s beam pointing obliquely.

In order to recover the electron density height profile, we assume a homogenous electron density profile over the EISCAT and KAIRA sites. This allows us to assume an electron density profile while taking into account



**Figure 8.** The height of the peak electron density and the peak electron density as output by IONONEST using (a) Model 1 and (b) Model 2 compared to the height of the peak electron density and the peak electron density measured by EISCAT.



**Figure 9.** (left column) The electron density at the points (A and B) at which the KAIRA beam intersect with the EISCAT beam. At an elevation of 45°, the KAIRA beam intersects with the EISCAT beam at point A, which is 85 km above the EISCAT site and 120 km along the KAIRA beam. At an elevation of 50°, the KAIRA beam intersects with the EISCAT beam at point B, which is 101 km above the EISCAT site and 132 km along the KAIRA beam. (right column) The electron density at the points 85 km and 101 km above the KAIRA site, taken from the Model 1 and Model 2 fits to the absorption data obtained by IONONEST from the KAIRA zenith pointing.

the increased absorption due to the longer path length. The electron density height profile is mapped to the path, and the values along the path length are interpolated. The absorption occurring along this path is then calculated in the same way as in sections 4 and 5.

Figure 9 shows comparisons of the EISCAT electron density at the points where the beams intersect with the electron densities as calculated by IONONEST from the KAIRA absorption data for the pointings over the EISCAT VHF. For comparison, it also shows the electron densities at the same heights as calculated by IONONEST for the two models from the KAIRA absorption data for the zenith pointing. The electron densities obtained for the pointings above the EISCAT site are more similar to the EISCAT data than those obtained by the KAIRA beam pointing at zenith. The fit to pointing B appears better than that of pointing A when compared to the EISCAT data, though it still fails to recreate the peak in electron density at 04:00–05:00 UT. However, a peak in electron density is seen at 02:00–03:00 UT in the KAIRA data as well as the EISCAT data, although the EISCAT peak is stronger than the KAIRA peak. This could, again, be caused by the mismatch in beamsizes. Between 04:00 and 05:00 UT, the absorption measured for each frequency with KAIRA is more prominent in the high-frequency (and thus smaller) beams, implying that the absorption structure is not a broad, all-sky structure.

The difference between the EISCAT and KAIRA electron densities could also be caused by density enhancements not contributing to the absorption with the same effectiveness at heights of 100 km because, although the electron density becomes high, the neutral density is relatively low, and hence, the electron-neutral collision frequency is lower at this altitude (as shown in Figure 1). The majority of the electron density at



04:00–05:00 UT is located at a higher altitude; therefore, it contributes less to the absorption measurements made by KAIRA, which could lead to an underestimation in the electron density.

To investigate the true causes of the peaks in electron density between 02:00–03:00 UT and 04:00–05:00 UT, all-sky observations corresponding to the EISCAT and KAIRA observations could be used, allowing for comparison with auroral activity. There is no increase in  $AE$  index at 04:00–05:00 UT, implying that the peak in electron density could be due to an increase in localized auroral activity which would be seen by all-sky observations. Unfortunately, these could not be obtained because although there are all-sky cameras located at the EISCAT site in Ramfjordmoen, there was cloud cover on the day these observations were taken.

The quality of fit is again assessed using the SMSE (equation (10)), setting  $Z_i$  to be the KAIRA absorption data for the two pointing directions,  $Z_{var}$  to be the variance of the KAIRA absorption data for each pointing direction and  $Z_i^*$  to be the Model 1 or Model 2 fit to the KAIRA absorption data. The average SMSE across the frequency range for Model 1 is found to be 0.11 for pointing A and 0.13 for B, while for Model 2 it is 0.13 for A and 0.15 for B, showing that both models are good fits to the KAIRA absorption data.

## 6. Discussion and Conclusions

IONONEST, a new technique for obtaining electron density height profiles for the  $D$  region of the ionosphere from multifrequency absorption measurements, has been introduced. IONONEST enables us to find the posterior probability distribution of parameters that describes an electron density height profile by comparing measured or simulated absorption data to absorption data calculated from a parameterized electron density height profile model. IONONEST has the benefit over many other inverse methods of returning uncertainties on these returned parameters, as well as sampling efficiently from multimodal or degenerate posterior distributions, something that standard Markov Chain Monte Carlo methods, such as the Metropolis-Hastings algorithm, struggle with. IONONEST returns Bayesian evidence values for electron density height profile models, giving IONONEST the ability to compare different electron density height profile models.

We have demonstrated that IONONEST can return accurate parameters for ionospheric electron density height profiles from multifrequency absorption data calculated using simulated ionospheric electron density height profiles. The parameters for the electron density height profiles are returned within  $\pm 1\sigma$  of the input parameters for realistic noise conditions.

We have also used IONONEST to fit to KAIRA multifrequency absorption data and have compared the returned electron density height profiles to data from the EISCAT VHF to qualitatively verify the fits. This demonstrates that IONONEST is capable of returning realistic electron density height profiles that are comparable (within  $3\sigma$ ) to those obtained by the EISCAT VHF, bearing in mind the differences between the two instruments. The electron density height profiles returned using IONONEST for the KAIRA pointings over the EISCAT VHF site showed similar temporal variability and magnitude to the measured EISCAT VHF electron densities at those heights.

Unlike EISCAT, KAIRA is capable of providing continuous measurements of absorption which can then be converted to continuous electron density height profiles. While EISCAT has the advantage of being able to measure electron densities at higher altitudes, KAIRA has the potential to return electron density height profiles that extend down to lower altitudes. This can be achieved by using IONONEST in conjunction with an ion chemistry model that extends to lower altitudes than the NRL-MSISE-00 reference atmosphere which is used in this paper to calculate the electron-neutral collision frequency.

In this paper we have considered two  $D$  region models for ionospheric electron density height profiles, a two parameter exponential model (Model 1) and a polynomial model (Model 2). By comparing Bayes factors and SMSE values, we have found that Model 2 is capable of recreating more realistic electron density height profiles, while Model 1 consistently overestimates the electron density at the top of the modeled height range.

### 6.1. Extensions to the Method

Having shown that IONONEST is capable of recreating electron density height profiles from absorption measurements from the KAIRA instrument, this method could now be extended to similar facilities such as LOFAR and the LWA, as well as future radio telescopes operating at low frequencies such as the low-frequency Square Kilometre Array (SKA-Low) [Garrett *et al.*, 2010]. SKA-Low is, however, limited in its ability to measure absorption as its minimum frequency is 50 MHz. The use of low-frequency radio telescopes could enable continuous measurements of  $D$  region electron density height profiles over Europe, Australia, and South Africa, allowing

for observations to be made simultaneously at a range of midlatitudes. Models such as the International Reference Ionosphere (IRI) [Bilitza *et al.*, 2011] could use these continuous measurements to improve their models of the currently underobserved *D* region.

In this paper we have compared just two electron density height profile models; however, there are many more within the literature that could be implemented in IONONEST. Furthermore, these data provide a powerful resource for developing *D* region models further. Monitoring observations also raises the possibility of probing *D* region changes on short time scales, for example, during solar events.

#### Acknowledgments

KAIRA was funded by the University of Oulu and the FP7 European Regional Development Fund and is operated by Sodankylä Geophysical Observatory. EISCAT is an international association supported by research organizations in China (CRIRP), Finland (SA), Japan (NIPR and STEL), Norway (NFR), Sweden (VR), and the United Kingdom (NERC). We gratefully acknowledge the support from NERC for access to the EISCAT facility. P.L.M. gratefully acknowledges support from the South East Physics Network (SEPnet). A.M.S. gratefully acknowledges support from the European Research Council under grant ERC-2012-StG-307215 LODESTONE. P.L.M. wishes to thank A. Kero for helpful discussions in the early stages of this paper. Following the data policy of AGU, all the data presented here are curated and made available upon request by the corresponding author.

#### References

- Aksnes, A., J. Stadsnes, N. Østgaard, G. A. Germany, K. Oksavik, R. R. Vondrak, A. Brekke, and U. P. Løvhaug (2006), Height profiles of the ionospheric electron density derived using space-based remote sensing of UV and X-ray emissions and EISCAT radar data: A ground-truth experiment, *J. Geophys. Res.*, *111*, A02301, doi:10.1029/2005JA011331.
- Baron, M. (1986), EISCAT progress 1983–1985, *J. Atmos. Terr. Phys.*, *48*, 767–772.
- Bilitza, D., L.-A. McKinnell, B. Reinisch, and T. Fuller-Rowell (2011), The international reference ionosphere today and in the future, *J. Geod.*, *85*, 909–920, doi:10.1007/s00190-010-0427-x.
- Browne, S., J. K. Hargreaves, and B. Honary (1995), An imaging riometer for ionospheric studies, *Electron. Commun. Eng. J.*, *7*(5), 209–217.
- Budden, K. G. (1985), *The Propagation of Radio Waves: The Theory of Radio Waves of Low Power in the Ionosphere and Magnetosphere*, Cambridge Univ. Press, Cambridge, U. K.
- Cameron, E., and A. Pettitt (2013), Recursive pathways to marginal likelihood estimation with prior-sensitivity analysis, ArXiv e-prints.
- Cheng, Z., S. A. Cummer, D. N. Baker, and S. G. Kanekal (2006), Nighttime *D* region electron density profiles and variabilities inferred from broadband measurements using VLF radio emissions from lightning, *J. Geophys. Res.*, *111*, A05302, doi:10.1029/2005JA011308.
- Detrick, D. L., and T. J. Rosenberg (1990), A phased-array radiowave imager for studies of cosmic noise absorption, *Radio Sci.*, *25*, 325–338, doi:10.1029/RS025i004p00325.
- Ellingson, S. W., T. E. Clarke, A. Cohen, J. Craig, N. E. Kassim, Y. Pihlstrom, L. J. Rickard, and G. B. Taylor (2009), The long wavelength array, *IEEE Proc.*, *97*, 1421–1430, doi:10.1109/JPROC.2009.2015683.
- Feroz, F., and M. P. Hobson (2008), Multimodal nested sampling: An efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses, *Mon. Not. R. Astron. Soc.*, *384*, 449–463, doi:10.1111/j.1365-2966.2007.12353.x.
- Feroz, F., P. J. Marshall, and M. P. Hobson (2008), Cluster detection in weak lensing surveys, ArXiv e-prints.
- Feroz, F., M. P. Hobson, and M. Bridges (2009), MULTINEST: An efficient and robust Bayesian inference tool for cosmology and particle physics, *Mon. Not. R. Astron. Soc.*, *398*, 1601–1614, doi:10.1111/j.1365-2966.2009.14548.x.
- Folkestad, K., T. Hagfors, and S. Westerlund (1983), EISCAT—An updated description of technical characteristics and operational capabilities, *Radio Sci.*, *18*, 867–879, doi:10.1029/RS018i006p00867.
- Friedrich, M., and M. Rapp (2009), News from the lower ionosphere: A review of recent developments, *Surv. Geophys.*, *30*, 525–559, doi:10.1007/s10712-009-9074-2.
- Garrett, M. A., J. M. Cordes, D. R. Deboer, J. L. Jonas, S. Rawlings, and R. T. Schilizzi (2010), Square Kilometre Array: A concept design for phase 1, ArXiv e-prints.
- Hargreaves, J. K. (1992), *The Solar-Terrestrial Environment: An Introduction to Geospace—The Science of the Terrestrial Upper Atmosphere, Ionosphere and Magnetosphere*, Cambridge Univ. Press, New York.
- Hargreaves, J. K., and M. Friedrich (2003), The estimation of *D*-region electron densities from riometer data, *Ann. Geophys.*, *21*, 603–613, doi:10.5194/angeo-21-603-2003.
- Honary, F., S. R. Marple, K. Barratt, P. Chapman, M. Grill, and E. Nielsen (2011), Invited article: Digital beam-forming imaging riometer systems, *Rev. Sci. Instrum.*, *82*(3), 031301.
- Hunsucker, R. D., and J. K. Hargreaves (2002), *The High-Latitude Ionosphere and Its Effects on Radio Propagation*, Cambridge Univ. Press, Cambridge, U. K.
- Jeffreys, H. (1961), *Theory of Probability*, 3rd ed., Oxford Univ. Press, New York.
- Kero, A., J. Vierinen, D. McKay-Bukowski, C.-F. Enell, M. Sinor, L. Roininen, and Y. Ogawa (2014), Ionospheric electron density profiles inverted from a spectral riometer measurement, *Geophys. Res. Lett.*, *41*, 5370–5375, doi:10.1002/2014GL060986.
- Lentati, L., P. Alexander, M. P. Hobson, F. Feroz, R. van Haasteren, K. J. Lee, and R. M. Shannon (2014), TEMPONEST: A Bayesian approach to pulsar timing analysis, *Mon. Not. R. Astron. Soc.*, *437*, 3004–3023, doi:10.1093/mnras/stt2122.
- Little, C., and H. Leinbach (1959), The riometer—A device for the continuous measurement of ionospheric absorption, *Proc. IRE*, *47*(2), 315–320, doi:10.1109/JRPROC.1959.287299.
- McKay-Bukowski, D., *et al.* (2015), Kaira: The Kilpisjärvi atmospheric imaging receiver array—System overview and first results, *IEEE Trans. Geosci. Remote Sens.*, *53*(3), 1440–1451, doi:10.1109/TGRS.2014.2342252.
- Parthasarathy, R., G. M. Lerfeld, and C. G. Little (1963), Derivation of electron-density profiles in the lower ionosphere using radio absorption measurements at multiple frequencies, *J. Geophys. Res.*, *68*, 3581–3588, doi:10.1029/JZ068i012p03581.
- Picone, J. M., A. E. Hedin, D. P. Drob, and A. C. Aikin (2002), NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues, *J. Geophys. Res.*, *107*(A12), 1468, doi:10.1029/2002JA009430.
- Rasmussen, C. E., and C. K. I. Williams (2005), *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, MIT Press, Cambridge, U. K.
- Rishbeth, H., and A. P. van Eyken (1993), EISCAT—Early history and the first ten years of operation, *J. Atmos. Terr. Phys.*, *55*, 525–542.
- Sen, H. K., and A. A. Wyller (1960), On the generalization of the Appleton-Hartree magnetoionic formulas, *J. Geophys. Res.*, *65*, 3931–3950, doi:10.1029/JZ065i012p03931.
- Skilling, J. (2004), Nested Sampling, in *24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, *Am. Inst. of Phys. Conf. Ser.*, vol. 735, edited by R. Fischer, R. Preuss, and U. V. Toussaint, pp. 395–405, Am. Inst. of Phys., Melville, New York, doi:10.1063/1.1835238.
- Sokolowski, M., R. B. Wayth, S. E. Tremblay, S. J. Tingay, M. Waterson, J. Tickner, D. Emrich, F. Schlagenhauer, D. Kenney, and S. Padhi (2015), The impact of the ionosphere on ground-based detection of the global epoch of reionization signal, *Astrophys. J.*, *813*, 18, doi:10.1088/0004-637X/813/1/18.

- Stauning, P. (1996), Investigations of ionospheric radio wave absorption processes using imaging riometer techniques, *J. Atmos. Terr. Phys.*, *58*, 753–764, doi:10.1016/0021-9169(95)00072-0.
- Stoker, P. H. (1987), Riometer absorption and spectral index for precipitating electrons with exponential spectra, *J. Geophys. Res.*, *92*, 5961–5968, doi:10.1029/JA092iA06p05961.
- Thompson, A. R., J. M. Moran, and G. W. Swenson (1986), *Interferometry and Synthesis in Radio Astronomy*, Wiley, New York.
- Thomson, N. R. (1993), Experimental daytime VLF ionospheric parameters, *J. Atmos. Terr. Phys.*, *55*, 173–184.
- Trotta, R. (2008), Bayes in the sky: Bayesian inference and model selection in cosmology, *Contemp. Phys.*, *49*, 71–104, doi:10.1080/00107510802066753.
- van Haarlem, M. P., et al. (2013), LOFAR: The low-frequency array, *Astron. Astrophys.*, *556*, A2, doi:10.1051/0004-6361/201220873.
- Vedantham, H. K., L. V. E. Koopmans, A. G. de Bruyn, S. J. Wijnholds, B. Ciardi, and M. A. Brentjens (2014), Chromatic effects in the 21 cm global signal from the cosmic dawn, *Mon. Not. R. Astron. Soc.*, *437*, 1056–1069, doi:10.1093/mnras/stt1878.
- Vierinen, J., D. McKay-Bukowski, M. Lehtinen, A. Kero, and T. Ulich (2013), Kilpisjärvi Atmospheric Imaging Receiver Array—First results, paper presented at 2013 IEEE International Symposium on Phased Array Systems Technology, pp. 664–668, doi:10.1109/ARRAY.2013.6731909.
- Wait, J., and K. Spies (1964), Propagation of radio waves past a coast line with a gradual change of surface impedance, *IEEE Trans. Antennas Propag.*, *12*, 570–575, doi:10.1109/TAP.1964.1138273.