AGU PUBLICATIONS

Radio Science

RESEARCH ARTICLE
10.1002/2014RS005431

Special Section:
Beacon Satellite Symposium 2013

This article is a companion to *Vierinen et al.* [2014] doi:10.1002/2014RS005434.

Key Points:
- We present a novel ionospheric tomography reconstruction method
- The method is based on Bayesian inference with the use of GMRF priors
- The prior distribution is built based on physical and empirical information

Correspondence to:
J. Norberg,
johannes.norberg@fmi.fi

Citation:
Norberg, J., L. Roininen, J. Vierinen, O. Amm, D. McKay-Bukowski, and M. Lehtinen (2015), Ionospheric tomography in Bayesian framework with Gaussian Markov random field priors, *Radio Sci.*, *50*, 138–152, doi:10.1002/2014RS005431.

Received 14 MAR 2014
Accepted 21 JAN 2015
Accepted article online 28 JAN 2015
Published online 18 FEB 2015

Ionospheric tomography in Bayesian framework with Gaussian Markov random field priors

J. Norberg[1], L. Roininen[2], J. Vierinen[3], O. Amm[1], D. McKay-Bukowski[2], and M. Lehtinen[2]

[1]Finnish Meteorological Institute, Helsinki, Finland, [2]Sodankylä Geophysical Observatory, University of Oulu, Sodankylä, Finland, [3]Haystack Observatory, Massachusetts Institute of Technology, Westford, Massachusetts, USA

Abstract We present a novel ionospheric tomography reconstruction method. The method is based on Bayesian inference with the use of Gaussian Markov random field priors. We construct the priors as a system of stochastic partial differential equations. Numerical approximations of these equations can be represented with linear systems with sparse matrices, therefore providing computational efficiency. The method enables an interpretable scheme to build the prior distribution based on physical and empirical information on the structure of the ionosphere. We show through synthetic test cases in a two-dimensional setup of latitude-altitude slices how this method can be applied to satellite-based ionospheric tomography and how information about the structure of the ionosphere can be implemented in the prior. The technique is capable of being easily extended to multifrequency tomographic analysis or used for the inclusion of other data sets of ionospheric electron density, such as ground-based observations by radars or ionosondes.

1. Introduction

The object of ionospheric tomography is to study the electron density of Earth's upper atmosphere between altitudes of 60 and 1000 km. The idea was originally proposed by *Austen et al.* [1988], and the statistical approach for the problem was first presented by *Fremouw et al.* [1992]. The first to utilize Bayesian statistical theory and Gaussian prior information for ionospheric tomography was *Markkanen et al.* [1995]. Later studies in the field are concentrated mostly in the three-dimensional cases and real-time estimations [*Hansen*, 2002; *Mitchell and Spencer*, 2003] where the problem is essentially the same as that in two dimensions, but the sparseness of the data and increasing computational demand become even more crucial factors. More recently, *Seemala et al.* [2014] have used algorithm similar to *Markkanen et al.* [1995] for three-dimensional ionospheric tomography. The method development of ionospheric tomography is well documented by *Bust and Mitchell* [2008]. Standard ionospheric tomography techniques are covered in text books, such as *Kunitsyn and Tereshchenko* [2003].

The main data component in ionospheric tomography is usually provided with ground-based measurements of signals transmitted by Global Positioning System (GPS) satellites or Low Earth Orbit (LEO) beacon satellites. Here we concentrate on a simple case that models the overflight of one LEO satellite and is illustrated in Figure 1. However, the method presented here can utilize various types of ionospheric measurements.

The unknown object of interest is the electron density in the ionosphere. We assume that it is constant in time during the flyby, which for an overflight of a LEO satellite over a regional tomographic receiver network on the ground, like the TomoScand receivers in Fenno-Scandinavia, is approximately 15 min *Vierinen et al.* [2014].

Mathematically, ionospheric tomography is an ill-posed problem which can be considered a sparse limited-angle tomography and an inverse problem [*Kaipio and Somersalo*, 2005]. The low information provided by the measurement setup in ionospheric tomography has been discussed in several articles. The restrictions caused by the limited elevation angle and limited number of receivers have been studied by *Yeh and Raymund* [1991], and the connection between measurement geometry and possible lattice resolutions are analyzed by *Na and Sutton* [1994]. *Saksman et al.* [1997] show that the ionosphere can contain structures that remain completely invisible for a certain class of receiver geometries.

As the information provided by the receiver-satellite geometry is very limited, the unknown electron density cannot be solved without stabilizing/regularizing the problem with some additional information.

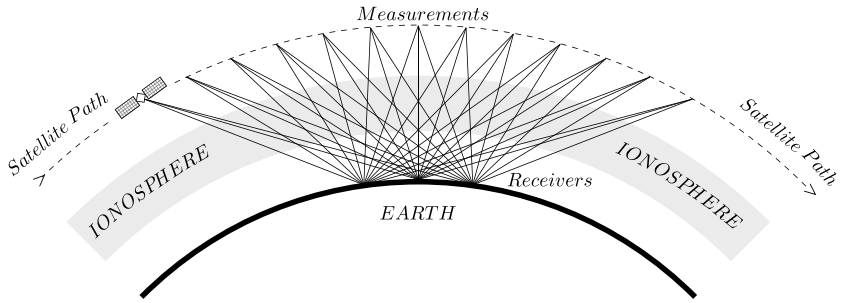

Figure 1. Schematic plot of ionospheric tomography.

The most commonly used methods in ionospheric tomography include iterative methods such as the Algebraic Reconstruction Technique [*Andreeva*, 1990] and similar iterative techniques [*Raymund et al.*, 1990; *Pryse et al.*, 1993], use of base functions in the vertical direction [*Na and Lee*, 1990; *Fremouw et al.*, 1992], and smoothness constraints [*Markkanen et al.*, 1995].

Raymund [1995] compared different algorithms and their variations for ionospheric tomography. While there were significant differences between the methods, it was also reported that the well performing algorithms did not have consistency over different cases. As one calibration of the model can produce good solution with some data, it can be very unsuitable under different conditions.

As the ionospheric tomography reconstructions are dominated by the additional information of the algorithm, it is important to understand what this information is and how strongly it affects the solution. In the Bayesian statistical approach, the stabilization for the inverse problem is given as an *a priori probability distribution*. The a priori probability distribution contains the information on the unknown parameters one has before actual measurements. The solution, i.e., the *a posteriori probability distribution* is constructed from the a priori distribution and the likelihood density functions. The latter, given the measurements, states the likelihood of unknown parameters. For the remainder of the paper we refer to these probability distributions simply as *prior* and *posterior* and to parameters of these distributions with the premodifiers prior and posterior.

Markkanen et al. [1995] used Bayesian approach for ionospheric tomography where the stabilization was provided by constraining the differences between neighboring pixels with an altitude-dependent regularization parameter profile. This approach gives a computationally efficient way to stabilize the problem by requiring smoothness from the reconstruction. However, this kind of smoothness constraint does not result in a proper prior probability distribution; i.e., the covariance matrix does not exist, and the regularizing effect is controlled with a regularization parameter that has no clear physical interpretation.

Natterer [1986] considered general computerized tomography where an isotropic exponential model can be understood as a prior covariance function for a Radon transform problem. For ionospheric tomography, *Arikan et al.* [2007] used similar approach with Gaussian random field priors with proper covariance matrices. This method is a truly Bayesian approach for ionospheric tomography, where the prior information can be given as a proper distribution. Instead of hiding the stabilization inside the algorithm, it is given with a clear physical interpretation. In practice, the prior distribution can be formed by choosing the prior mean and covariance structure based on, e.g., International Reference Ionosphere model, Chapman profile, other ionospheric models, or additional measurements such as profiles from ionosonde measurements. The variance of the system should then reflect the credibility of the given parameters.

Despite the evident advances of the Bayesian approach, the shortcoming, that was remarked also by *Arikan et al.* [2007], is that due to the resulting full covariance matrix, the problem becomes easily computationally unfeasible. This is the case especially in the three-dimensional ionospheric tomography where the number of unknown parameters rises rapidly.

Here we follow *Roininen et al.* [2011, 2013, 2014] to utilize similar smoothness constraints as in *Markkanen et al.* [1995] but parametrized in a more advanced manner so that the given constraints form a Gaussian Markov random field (GMRF). The GMRFs are implemented with sparse matrices providing computational efficiency, yet the approach provides a known and controllable covariance structure. Hence, the prior

probability distribution can be defined by prior mean and covariance, which are given as physical quantities in units of electrons per volume and length.

The GMRF prior provides also discretization invariancy, i.e., the information given with the prior, and hence, the resulting reconstructions are essentially not affected by the change of the lattice used for the modeling. However, most importantly, the GMRFs provide a computationally efficient way for truly Bayesian statistical inversion with proper prior distributions. To our knowledge, this is the first time that GMRF priors are used to address the problem of ionospheric tomography.

The paper is organized as follows: in section 2, we review the basic concepts of the Bayesian statistical inversion. The construction of prior distribution with GMRF is explained in section 3, and the altitude information and the corresponding parametric fields for GMRFs are covered in section 4. In section 5 we show how the Bayesian statistical inversion with the GMRF priors is applied to ionospheric tomography in practice. Following section 6, which has our conclusions and discussion, we briefly consider the future plans on the topic in section 7. The first results with real measurement data are presented in the companion paper *Vierinen et al.* [2014].

2. Bayesian Statistical Inversion

Figure 1 shows a schematic plot of ionospheric tomography. The dual-frequency signal is transmitted from the beacon satellites and received on the ground. The relative difference of ionospheric propagation delays of the two different frequencies can be measured tracking the phase of the two carriers. When subtracting one phase curve from a scaled version of the other, the contribution of tropospheric refraction, clock drifts, and uncertainties in the satellite position to measurements are canceled out and the relative electron content along the signal path remains, up to an unknown constant. The actual measurement can then be modeled as a line integral of electron density

$$m_t = \gamma_i + \int_{L_t} N_{\mathrm{e}}(z)\mathrm{d}z_t, \tag{1}$$

where m_t is the measured relative total electron content, t is the time, $N_e(z)$ the electron density of the domain, and $z := (z_1, z_2) \in \mathbb{R}^2$ the coordinate. The integral is defined over the measurement signal path L_t with coordinates z_t. The unknown constant γ_i is due to the 2π ambiguity of the initial phase difference and the index i refers to a certain receiver-satellite combination. For details see *Davies* [1965]. We assume a ray-casting approximation for the paths, i.e., the path from satellite to receiver is always a straight line. We call equation (1) the continuous observation model.

When the electron density and phase constants are known, we call the solution for m_t of equation (1) the forward problem. The inverse problem is then to solve the electron density $N_{\mathrm{e}}(z)$ and phase constants γ_i, when other parameters of the forward model are known.

For computations, we need a discrete approximation of the continuous observation model (1). We set $z_{j,k} := (z_{1,j}, z_{2,k}) := (jh_1, kh_2)$, where h_1 and h_2 are the discretization lengths, i.e., the horizontal and vertical dimensions of the pixel in the discretized two-dimensional domain, row index $j = 1, 2, \ldots, n_{\mathrm{row}}$ and column index $k = 1, 2, \ldots, n_{\mathrm{col}}$. It is well known that a discretization of a linear integral equation can be obtained by a discrete linear approximation. Let us write the approximation as a linear matrix equation

$$\mathbf{m} = \mathbf{A}\mathbf{x} \tag{2}$$

where $\mathbf{m} := (m_1, \ldots m_{n_{\mathrm{m}}})^T$ is the measurement vector containing the measured electron contents, $\mathbf{A}$ the theory matrix, and the vector $\mathbf{x} := (N_{\mathrm{e}}(z_{1,1}), \ldots, N_{\mathrm{e}}(z_{n_{\mathrm{row}},n_{\mathrm{col}}}), \gamma_1, \ldots, \gamma_n)^T$ contains the unknown electron densities and phase constants. The index n_m is the total number of measurements and n the total number of phase constants. The theory matrix $\mathbf{A}$ is formed according to the measurement geometry, and it approximates the line integrals given in equation (1). Here the matrix $\mathbf{A} \in \mathbb{R}^{(n_{\mathrm{row}}n_{\mathrm{col}})\times(n_{\mathrm{row}}n_{\mathrm{col}}+n_\gamma)}$ also includes the design matrix part which adds a correct γ_i parameter for the corresponding measurement.

In reality, the measurements are noisy. The noise can come from various sources, for example, radio frequency interference, thermal noise of the measurement device itself, and from small scale-changes in plasma that cannot be captured within the model resolution. For more detailed discussion on the error sources we refer to our companion paper *Vierinen et al.* [2014].

In statistical inverse problems also all the other variables are viewed as random variables. The randomness describes our degree of information concerning their realizations [*Kaipio and Somersalo*, 2005]. Hence, we change the paradigm from a deterministic matrix model of equation (2) to a stochastic matrix model

$$\mathbf{M} = \mathbf{AX} + \mathbf{e}, \tag{3}$$

where $\mathbf{e} = (e_1, \ldots, e_{n_m})^T$ is the measurement noise vector and all vectors **M**, **X**, and **e** are considered random variables. For the observed measurements we use **m** as a realization of *M*. We assume that the **e** is normally distributed with zero mean and $\mathbf{\Sigma}$ covariance. We denote this as $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$. We also assume that we know the covariance matrix $\mathbf{\Sigma}$, but we do not know realizations of the noise **e** pathwise. In practice the mean and covariance of the measurement error distribution are estimated from the data.

The solution to the Bayesian statistical inverse problem is the posterior probability distribution given with the Bayes' formula [*Kaipio and Somersalo*, 2005].

$$D_{\text{post}}(\mathbf{X}) = \frac{D(\mathbf{M}|\mathbf{X})D_{\text{pr}}(\mathbf{X})}{D(\mathbf{M})} \propto D(\mathbf{M}|\mathbf{X})D_{\text{pr}}(\mathbf{X}) \propto \exp\left(-\frac{1}{2}(\mathbf{M}-\mathbf{AX})^T\mathbf{\Sigma}^{-1}(\mathbf{M}-\mathbf{AX})\right) D_{\text{pr}}(\mathbf{X}),$$

where we use $\propto$ for proportionality, *D* is a Gaussian normal density function with subscript "post" for posterior and "pr" for prior. $D(\mathbf{M}|\mathbf{X})$ is a normal density function for a conditional distribution of **M** when **X** is given, and here it is obtained from the discrete approximation of equation (3). $D(\mathbf{M}|\mathbf{X})$ can also be understood as a likelihood function for realizations of parameter **X** given observed measurements $\mathbf{M} = \mathbf{m}$.

From now on, we assume that the continuous unknown electron density lies in $\mathbb{R}^2$. We further assume that the prior distribution of the discrete unknown electron density **X** is a Gaussian field, i.e., a multivariate normal distribution with a mean value $\boldsymbol{\mu}$ and a covariance $\mathbf{\Sigma}_{\text{pr}}$. Then the unnormalized posterior density is

$$D_{\text{post}}(\mathbf{X}) \propto \exp\left(-\frac{1}{2}\left((\mathbf{m}-\mathbf{AX})^T\mathbf{\Sigma}^{-1}(\mathbf{m}-\mathbf{AX}) + (\mathbf{X}-\boldsymbol{\mu})^T\mathbf{\Sigma}_{\text{pr}}^{-1}(\mathbf{X}-\boldsymbol{\mu})\right)\right),$$

from which we can calculate the maximum a posteriori (MAP) estimator, i.e., the most probable state of the ionosphere, given the prior information and measurements. We give the maximum a posteriori estimator by the matrix equation

$$\mathbf{x}_{\text{MAP}} = \mathbf{\Sigma}_{\text{post}}\left(\mathbf{A}^T\mathbf{\Sigma}^{-1}\mathbf{m} + \mathbf{\Sigma}_{\text{pr}}^{-1}\boldsymbol{\mu}\right), \tag{4}$$

where

$$\mathbf{\Sigma}_{\text{post}} = \left(\mathbf{A}^T\mathbf{\Sigma}^{-1}\mathbf{A} + \mathbf{\Sigma}_{\text{pr}}^{-1}\right)^{-1} \tag{5}$$

is the posterior covariance estimator. We note that the maximum a posteriori estimator is analytically the same as the generalized Tikhonov regularized solution but, in this case, carries a statistical interpretation [*Kaipio and Somersalo*, 2005].

The prior covariance matrix $\mathbf{\Sigma}_{\text{pr}}$ can be formed with a desired structure as in *Natterer* [1986] and *Arikan et al.* [2007]; however, the resulting covariance matrix $\mathbf{\Sigma}_{\text{pr}}$ is a full matrix, and therefore, the computation of the posterior covariance becomes computationally heavy.

Let us assume that we know a sparse matrix **L** which satisfies $\mathbf{L}^T\mathbf{L} = \mathbf{\Sigma}_{\text{pr}}^{-1}$. Then we do not need to invert the full covariance matrix. On the other hand, we can then represent the same information as a virtual measurement equation

$$\mathbf{L}(\mathbf{X}-\boldsymbol{\mu}) = \mathbf{W} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \tag{6}$$

where the vector **W** is normally distributed, the mean **0** a vector of zeros, and the covariance **I** an identity matrix. Then the model can be written as a system of matrix equations

$$\begin{bmatrix} \mathbf{M} \\ \mathbf{L}\boldsymbol{\mu} \end{bmatrix} = \begin{bmatrix} \mathbf{A} \\ \mathbf{L} \end{bmatrix} \mathbf{X} + \begin{bmatrix} \mathbf{e} \\ \mathbf{W} \end{bmatrix}. \tag{7}$$

It is then trivial to note that the weighted least squares solution of the system of equations (7) gives the maximum a posteriori estimate of equation (4). This is a critical point in the development of the method expounded in this paper, as we can present the full covariance matrix by a sparse matrix **L**. In the following two sections, we consider how to construct this matrix **L** as a sparse matrix.

3. Gaussian Markov Random Field Prior

The main goal of this article is to build the matrix $\mathbf{\Sigma}_{\text{pr}}^{-1}$ of equation (4) as a sparse matrix, but in such a way that it together with the prior mean $\boldsymbol{\mu}$ defines a proper normal distribution with known covariance structure. Here we do this with the formalism of the GMRFs.

The GMRF can be understood as a multivariate normal distribution, but instead of mean and covariance, the distribution is defined by its mean and precision matrix, i.e., the inverse covariance matrix $\mathbf{\Sigma}_{\text{pr}}^{-1}$. Here we concentrate on cases where the positive definite precision matrix is sparse. The zeros in the precision matrix indicate that the corresponding unknown is conditionally independent from these elements, given the nonzero entries. This is known as the Markov property. For a detailed definition of GMRFs we refer to *Rue and Held* [2005].

Here we follow *Roininen et al.* [2011, 2013] to build the matrix $\mathbf{L}$ of precision matrix $\mathbf{L}^T\mathbf{L} = \mathbf{\Sigma}_{\text{pr}}^{-1}$ in (4) so that the resulting GMRF covariance matrix approximates a continuous Gaussian-shaped covariance function. We first define and parametrize the distribution we are aiming at and then show how the corresponding matrix $\mathbf{L}$ is formed.

Let $X(z)$, $z \in \mathbb{R}^2$ be a real-valued Gaussian random function with finite moments of second order. We set

$$\begin{aligned} \mathbf{E}(X(z)) &= \mu(z) = \mu, \\ C(s) &:= \mathbf{E}(X(z)X(y)) = \mathbf{E}(X(z-y)X(0)), \end{aligned}$$

where $\mathbf{E}$ stands for the expectation operator and $s := (s_1, s_2) := z - y$, and $s, z, y \in \mathbb{R}^2$. Let us assume that $C(s)$ is an absolutely integrable function, then $X(z)$ is a weakly stationary random field. We then choose the prior distribution to be a weakly stationary Gaussian random field with mean μ and a squared exponential autocovariance function

$$C(s) = \alpha \exp\left(-\left(\frac{s_1^2}{l_1^2} + \frac{s_2^2}{l_2^2}\right)\right), \tag{8}$$

We call parameter α the scaling factor of the variance, and l_1, l_2 correlation lengths to respective coordinate directions.

We then present how the corresponding GMRF is constructed. We first define the discretized white noise as a discrete approximation of the continuous white noise W at a lattice point (j, k)

$$W_{j,k} \sim \mathcal{N}\left(0, \frac{1}{h_1 h_2}\right), \tag{9}$$

where h_1 and h_2 are the discretization lengths, i.e., the horizontal and vertical pixel dimensions of an uniform lattice. We denote the finite approximation of the continuous mean μ by vector $\boldsymbol{\mu} = (\bar{\mu}, \ldots, \bar{\mu})$. Then, according to *Roininen et al.* [2011, 2013], we can approximate the continuous spatial differential operator equations as a system of stochastic partial difference equations

$$\begin{cases} X_{j,k} - \bar{\mu} = \sqrt{\frac{\alpha l_1 l_2}{c_0 h_1 h_2}} W_{j,k}^{(0,0)} \\ X_{j,k} - X_{j-1,k} = \sqrt{\frac{\alpha h_1 l_2}{c_1 l_1 h_2}} W_{j,k}^{(1,1)} \\ X_{j,k} - X_{j,k-1} = \sqrt{\frac{\alpha h_1 l_2}{c_1 l_1 h_2}} W_{j,k}^{(1,2)} \\ \frac{l_1^2}{h_1^2}\left(X_{j+1,k} - 2X_{j,k} + X_{j-1,k}\right) + \\ \frac{l_2^2}{h_2^2}\left(X_{j,k+1} - 2X_{j,k} + X_{j,k-1}\right) = \sqrt{\frac{\alpha l_1 l_2}{c_2 h_1 h_2}} W_{j,k}^{(2,0)}, \end{cases} \tag{10}$$

where $W^{(p,p')}$ are independent discrete white noise processes as in (9), where the $p \in \{0, 1, 2, \ldots\}$ gives the order of differences and $p' \in \{0, 1, 2\}$ indicates the direction of differences with correspondence 1 = horizontal, 2 = vertical, and 0 = both directions; e.g., $W^{(1,2)}$ is a white noise process corresponding to first-order differences in vertical direction. Parameters c_p are known scaling constants chosen as explained in *Roininen et al.* [2013].

In *Roininen et al.* [2011, 2013] the equations in (10) were defined on whole space $\mathbb{R}$ or $\mathbb{R}^2$. For computations we change the paradigm from the whole space to a bounded domain. This is done via the introduction of periodic boundary conditions

$$X_{j,k} = X_{j+n_{\text{row}},k}, \text{ and } X_{j,k} = X_{j,k+n_{\text{col}}}. \tag{11}$$

Finally, we can write the discrete approximation in equation (10) as a stacked matrix equation corresponding to equation (6) as

$$\begin{bmatrix} \mathbf{L}^{(0,0)} \\ \mathbf{L}^{(1,1)} \\ \mathbf{L}^{(1,2)} \\ \mathbf{L}^{(2,0)} \end{bmatrix} (\mathbf{X} - \boldsymbol{\mu}) = \mathbf{W} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \tag{12}$$

where the matrices $\mathbf{L}^{(p,p')} \in \mathbb{R}^{n_{\text{row}}n_{\text{col}} \times n_{\text{row}}n_{\text{col}}}$ are defined as follows: The white noise matrix

$$\mathbf{L}^{(0,0)} = \sqrt{\frac{c_0 h_1 h_2}{\alpha l_1 l_2}}\mathbf{I}.$$

Let us denote by $\otimes$ the Kronecker product and an identity matrix by $\mathbf{I}_{n_{\text{row}} \times n_{\text{row}}} \in \mathbb{R}^{n_{\text{row}} \times n_{\text{row}}}$. Then the first-order difference matrices are

$$\mathbf{L}^{(1,1)} = \sqrt{\frac{c_1 l_1 h_2}{\alpha h_1 l_2}}\mathbf{I}_{n_{\text{row}} \times n_{\text{row}}} \otimes \begin{bmatrix} -1 & 1 & & \\ & \ddots & \ddots & \\ & & -1 & 1 \\ 1 & & & -1 \end{bmatrix}_{n_{\text{col}} \times n_{\text{col}}}$$

and

$$\mathbf{L}^{(1,2)} = \sqrt{\frac{c_1 h_1 l_2}{\alpha l_1 h_2}} \begin{bmatrix} -1 & 1 & & \\ & \ddots & \ddots & \\ & & -1 & 1 \\ 1 & & & -1 \end{bmatrix}_{n_{\text{row}} \times n_{\text{row}}} \otimes \mathbf{I}_{n_{\text{col}} \times n_{\text{col}}}.$$

The second-order differences can be presented as

$$\mathbf{L}^{(2,0)} = \sqrt{\frac{c_2 h_1 h_2}{\alpha l_1 l_2}} \left(\mathbf{I}_{n_{\text{row}} \times n_{\text{row}}} \otimes \begin{bmatrix} -2 & 1 & & & 1 \\ 1 & -2 & 1 & & \\ & \ddots & & \ddots & \\ & & 1 & -2 & 1 \\ 1 & & & 1 & -2 \end{bmatrix}_{n_{\text{col}} \times n_{\text{col}}} + \begin{bmatrix} -2 & 1 & & & 1 \\ 1 & -2 & 1 & & \\ & \ddots & & \ddots & \\ & & 1 & -2 & 1 \\ 1 & & & 1 & -2 \end{bmatrix}_{n_{\text{row}} \times n_{\text{row}}} \otimes \mathbf{I}_{n_{\text{col}} \times n_{\text{col}}} \right).$$

In *Roininen et al.* [2013] it is shown that by taking the order of difference matrices and the discretization interval into the limit, i.e., $p \to \infty$ and $h_1, h_2 \to 0$, the resulting covariance converges toward the continuous Gaussian-shaped covariance function given in equation (8). However, it is numerically verified that the matrices $\mathbf{L}^{(p,p')}$ can be used to form an adequate discrete approximation for this covariance already with $\max(p) = 2$, when h_1 and h_2 are shorter enough than the desired correlation lengths of the field.

As the discretization length is taken into account, the prior is discretization invariant and essentially the same information can be used for different lattice sizes. Also, the construction of priors on irregular lattices can be done similarly to the method used in this section. Any regular finite-difference approximation can be used in these cases. For the sake of notational clarity, we will not go through the construction of these priors on irregular lattices.

In *Roininen et al.* [2014] a similar idea was applied to a nonlinear Bayesian inverse problem, the electrical impedance tomography, but with finite element methods.

4. Altitude Information for the GMRF Prior

The electron density of the ionosphere varies spatially. For that, we change the constant parameters μ, α, l_1, l_2 introduced in the previous section to continuous, differentiable, and positive functions $\mu(z), \alpha(z), l_1(z), l_2(z)$ for locations z, where $z \in \mathbb{R}^2$.

Then the discrete GMRF approximation is

$$\begin{cases} X_{j,k} - \mu_{j,k} = \sqrt{\frac{\alpha_{j,k} l_{1,(j,k)} l_{2,(j,k)}}{c_0 h_1 h_2}} W_{j,k}^{(0)} \\ (X_{j,k} - \mu_{j,k}) - (X_{j-1,k} - \mu_{j-1,k}) = \sqrt{\frac{\alpha_{j,k} h_1 l_{2,(j,k)}}{c_1 l_{1,(j,k)} h_2}} W_{j,k}^{(1,1)} \\ (X_{j,k} - \mu_{j,k}) - (X_{j,k-1} - \mu_{j,k-1}) = \sqrt{\frac{\alpha_{j,k} l_{1,(j,k)} h_2}{c_1 h_1 l_{2,(j,k)}}} W_{j,k}^{(1,2)} \\ \frac{l_{1,(j,k)}^2}{h_1^2}((X_{j+1,k} - \mu_{j+1,k}) - 2(X_{j,k} - \mu_{j,k}) + (X_{j-1,k} - \mu_{j-1,k})) + \frac{l_{2,(j,k)}^2}{h_2^2}((X_{j,k+1} - \mu_{j,k+1}) - 2(X_{j,k} - \mu_{j,k}) \\ \quad +(X_{j,k-1} - \mu_{j,k-1})) = \sqrt{\frac{\alpha_{j,k} l_{1,(j,k)} l_{2,(j,k)}}{c_1 h_1 h_2}} W_{j,k}^{(2)}, \end{cases} \tag{13}$$

where the parameters $\mu_{j,k}, \alpha_{j,k}, l_{1(j,k)}$, and $l_{2(j,k)}$ are the values of the corresponding continuous functions at the location $z_{j,k}$. The approximation can then be presented with a similar stacked matrix representation as in equation (12).

In section 5, we show through numerical simulations that in practice the prior distribution can be determined with the parameters $\mu_{j,k}, \alpha_{j,k}, l_{1(j,k)}$, and $l_{2(j,k)}$ and that the effect of the parameters is interpretable. With the resulting GMRF prior, we can model structural properties of the ionosphere; i.e., we have an inhomogeneous prior. This approach is somewhat similar to the work of *Bardsley* [2013] and to *Kaipio et al.* [1999], in that we model the structural properties of the priors.

We note that from the purely mathematical point of view, we have not verified that there is a continuous GMRF prior in the discretization limit $h_1, h_2 \to 0$. However, numerical simulations support this idea, and we conjecture that the resulting GMRF prior is discretization invariant. Hence, the reconstructions computed are essentially independent of the computational lattice used.

5. Numerical Examples

In this section we show how the prior distributions are built with GMRFs and demonstrate with synthetical examples how the Bayesian approach with GMRF priors is implemented for ionospheric tomography. The first results with real data are presented in our companion paper *Vierinen et al.* [2014].

In the examples, we first simulate samples from the constructed prior distributions to show that the distributions actually incorporate the information given with the parameters of previous sections. We then show how well the unknown electron density, with a certain altitude characteristics and scale size of structures, can be reconstructed in an optimal case where we know the exact prior distribution. However, as with real measurements, the exact prior distribution is never actually known; we then show how well the unknown electron density can be reconstructed with a loose prior that does not restrict the solution with strict information that we actually do not have.

We take the theory presented in previous sections and build the inverted prior covariance matrix with matrix

$$\mathbf{L} = \begin{bmatrix} \mathbf{L}^{(0,0)} \\ \mathbf{L}^{(1,1)} \\ \mathbf{L}^{(1,2)} \\ \mathbf{L}^{(2,0)} \end{bmatrix},$$

where matrices $\mathbf{L}^{(0,0)}, \mathbf{L}^{(1,1)}, \mathbf{L}^{(1,2)}$, and $\mathbf{L}^{(2,0)}$ are constructed by following equation (13). In the vertical direction, we give the mean value $\boldsymbol{\mu} = (\mu_{1,1}, \mu_{1,2} \ldots, \mu_{n_{\text{row}}, n_{\text{col}}})^T$ and standard deviation $\sqrt{\boldsymbol{\alpha}} = \left(\sqrt{\alpha_{1,1}}, \sqrt{\alpha_{1,2}}, \ldots, \sqrt{\alpha_{n_{\text{row}}, n_{\text{col}}}}\right)^T$ of the field with a Chapman profile [*Brekke*, 1997].

Chapman profiles are altitude profiles of the ionospheric electron density, which result from the balance between electron production and loss processes in the ionosphere. Specifically, as a production process the

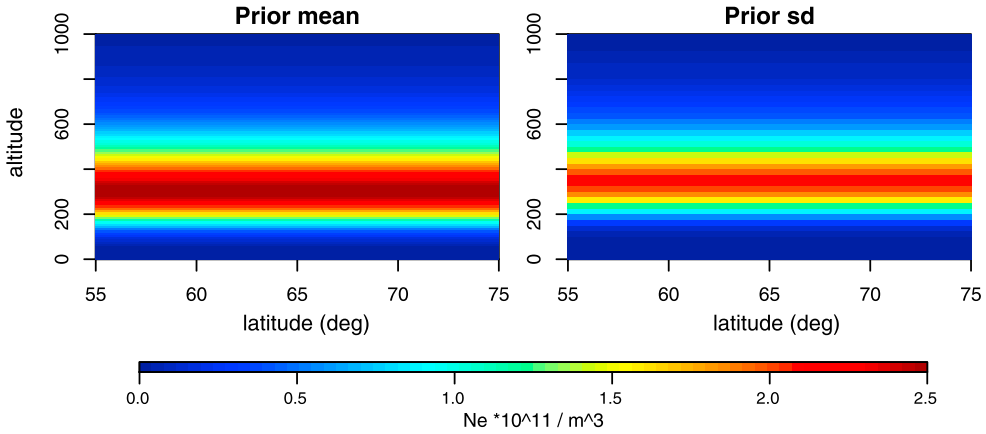

Figure 2. Parameter fields for prior distribution used in Prior 1.

Chapman profiles include electrochemical ionization due to energy input into the neutral atmosphere, and as a loss process recombination due to interaction of electrons with ions and neutral particles. The shape of the Chapman profile is defined with scale height, peak value, and peak altitude. The latter two of these parameters define what the expected maximum of the electron density is and its altitude. Scale height gives

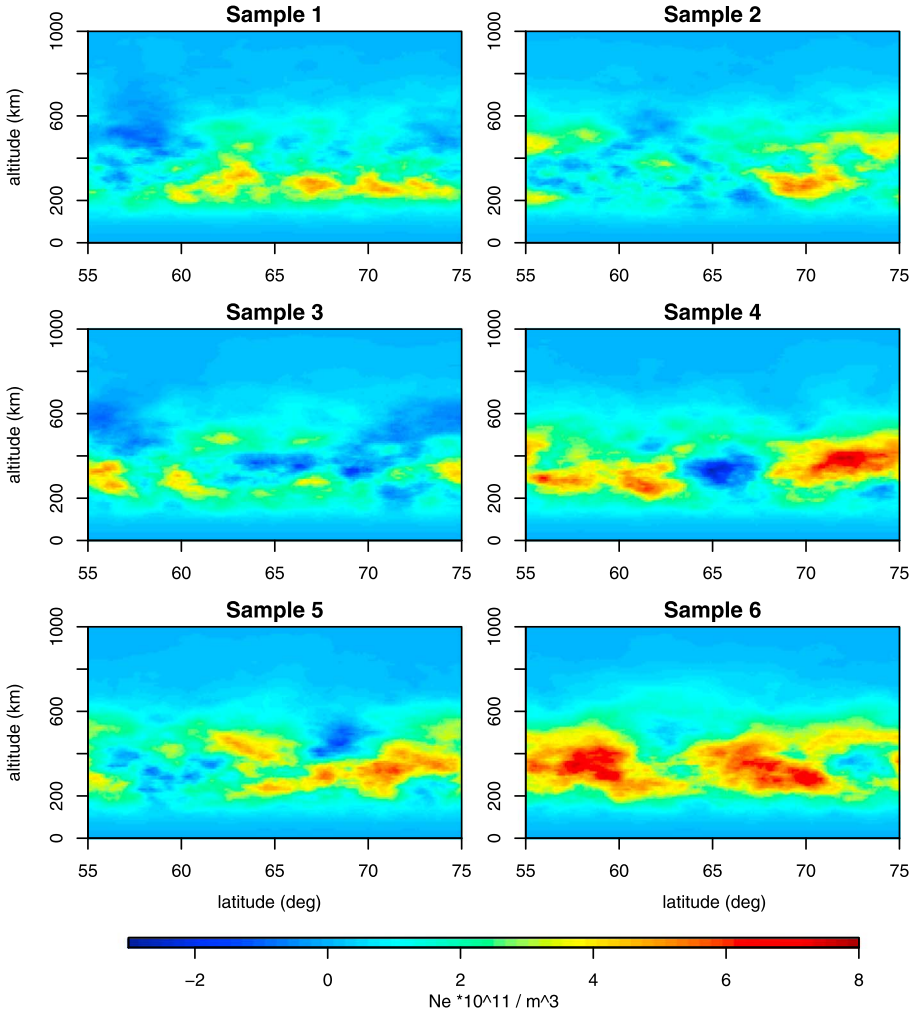

Figure 3. Six samples from the Prior 1 distribution.

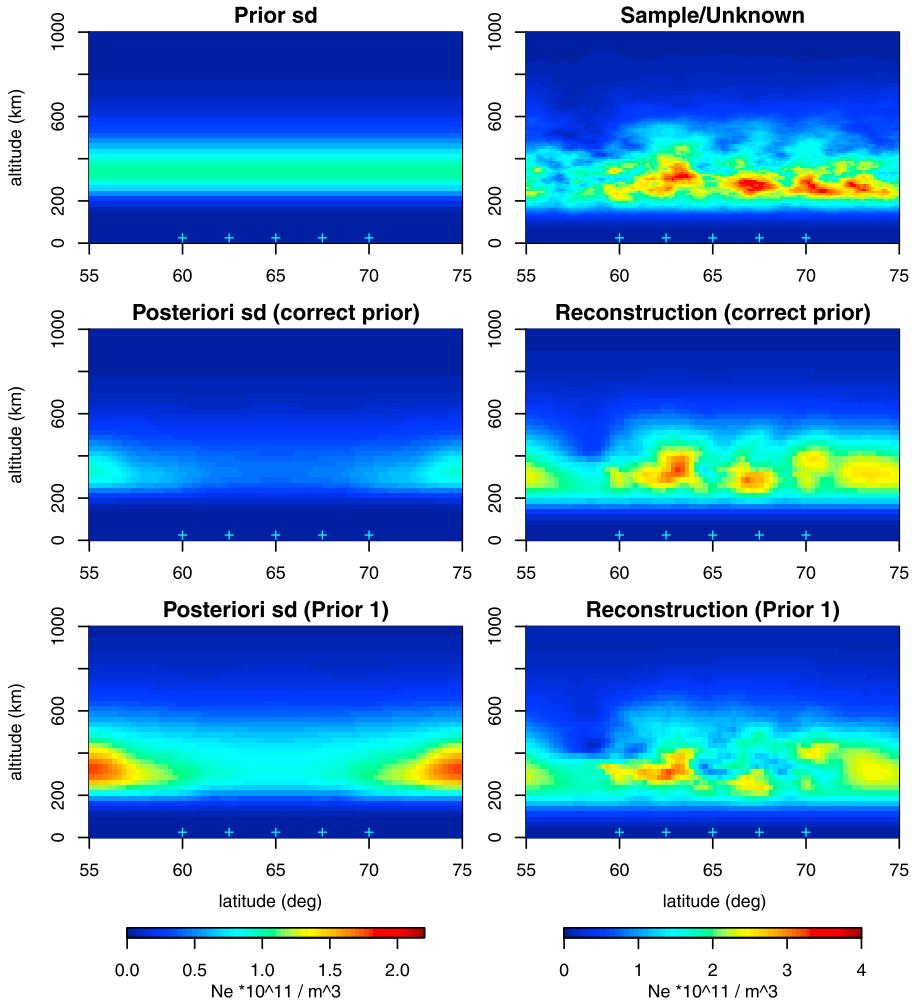

Figure 4. Case 1. (top row) The new prior standard deviation and the corresponding sample. (middle row) The posterior standard deviation and maximum a posteriori estimate with the correct prior. (bottom row) The posterior standard deviation and maximum a posteriori estimate with Prior 1. Receivers marked with cyan plus signs.

the vertical width of the profile. Shorter scale heights, close to 100 km or under, gives a profile where the electron density is expected to be concentrated strictly around the peak altitude. Larger scale heights such as 145 km used in the simulations here spreads the profile of expected electron density wider to low and high altitudes.

The mean value is our initial guess on the unknown electron density before the measurements, and the standard deviation states how much we trust this initial guess. In addition to that, we set a correlation length for both horizontal l_1 and vertical directions l_2. We define here the correlation length as a distance where the correlation between two points drops to 10% of variance. These parameters can be interpreted as how large details we expect to detect in the reconstruction. Very long correlation lengths correspond to a very smooth ionosphere.

The periodicity conditions of equation (11) are given for computational reasons. As these conditions cannot be justified physically, in the vertical direction we have cut the correlation length close to the vertical dimension of the pixel size at the boundaries. This could also be done for the horizontal direction, but in practice the reconstructions have to be done in domains larger than it is reasonable to use for analysis. Hence, the possible boundary effects are cut out from the resulting domain for ionospheric analysis.

In Figure 2 we present the mean value and standard deviations for an example case, from now on referred to as Prior 1. The peak for mean value is $2.5 \cdot 10^{11}\mathrm{Ne/m}^3$ at 300 km altitude. For the standard deviation the peak is $2 \cdot 10^{11}\mathrm{Ne/m}^3$, also at 300 km altitude. The scale height for the standard deviation profile is 120 km

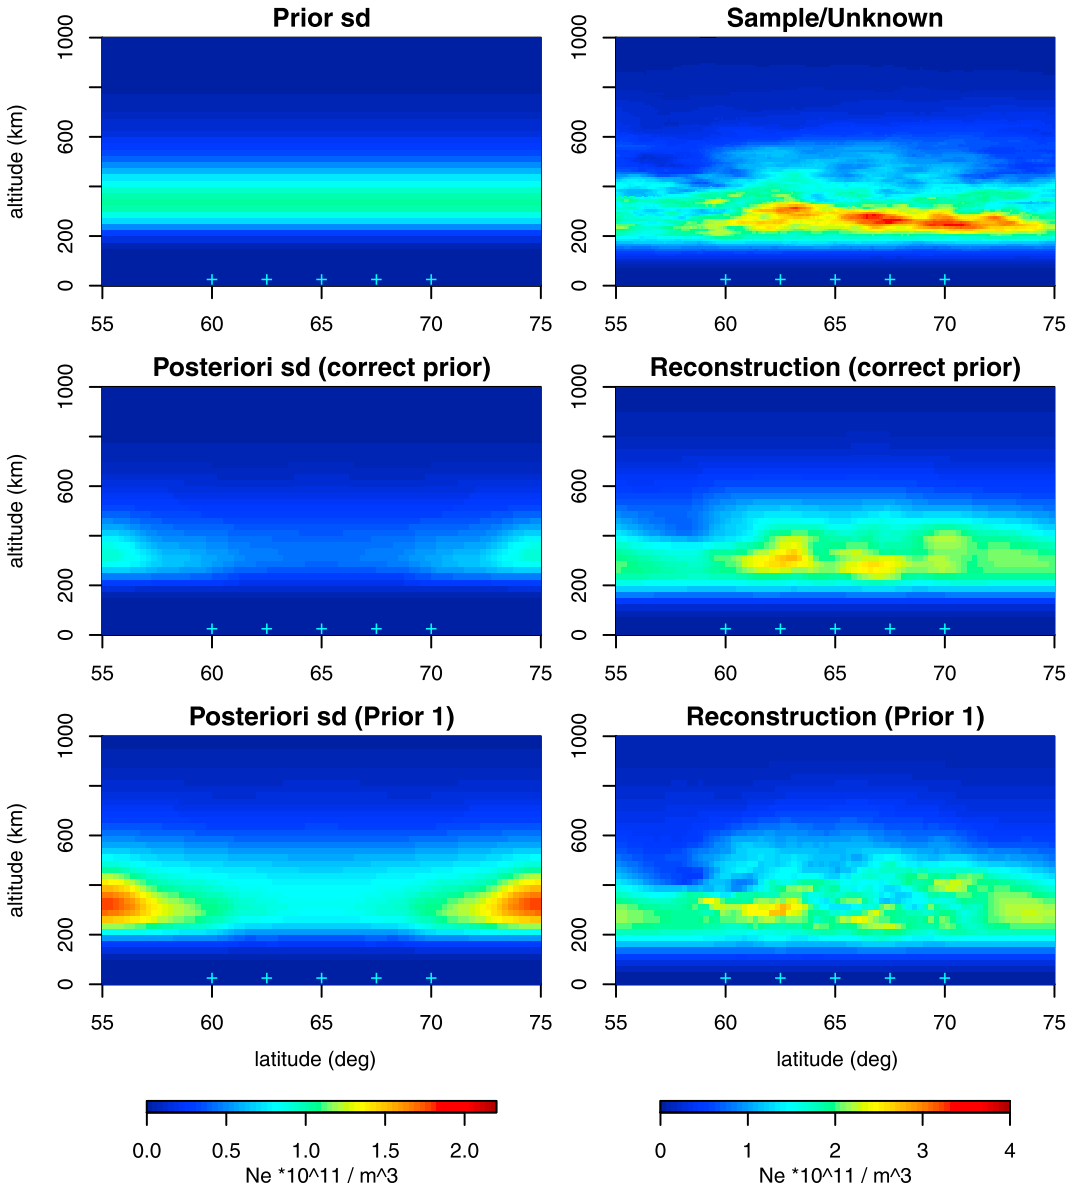

Figure 5. Case 2. (top row) The new prior standard deviation and the corresponding sample. (middle row) The posterior standard deviation and maximum a posteriori estimate with the correct prior. (bottom row) The posterior standard deviation and maximum a posteriori estimate with Prior 1. Receivers marked with cyan plus signs.

and 145 km for the mean. The horizontal correlation length is set to 10° and vertical to 400 km. In Figure 3 we have drawn six sample realizations from Prior 1 distribution. The samples are simulated by multiplying a standard normal random vector with Cholesky decomposition of the prior covariance matrix, then adding the prior mean value [*Gentle*, 2005].

The samples can be considered as likely realizations of the ionospheric electron density based only on our existing knowledge before the measurements. The samples provide us with a way to understand the prior distribution and to observe that the information stated in previous paragraphs is really incorporated in the samples. At this point we are not using any positivity constraints and, as the standard deviation of the prior is so large, there are some areas with negative values. We do not claim that these samples would be physically realistic realizations of the ionosphere but that we can clearly characterize structural information with the distributions as the samples clearly correspond to the chosen parameter values. The idea here is to

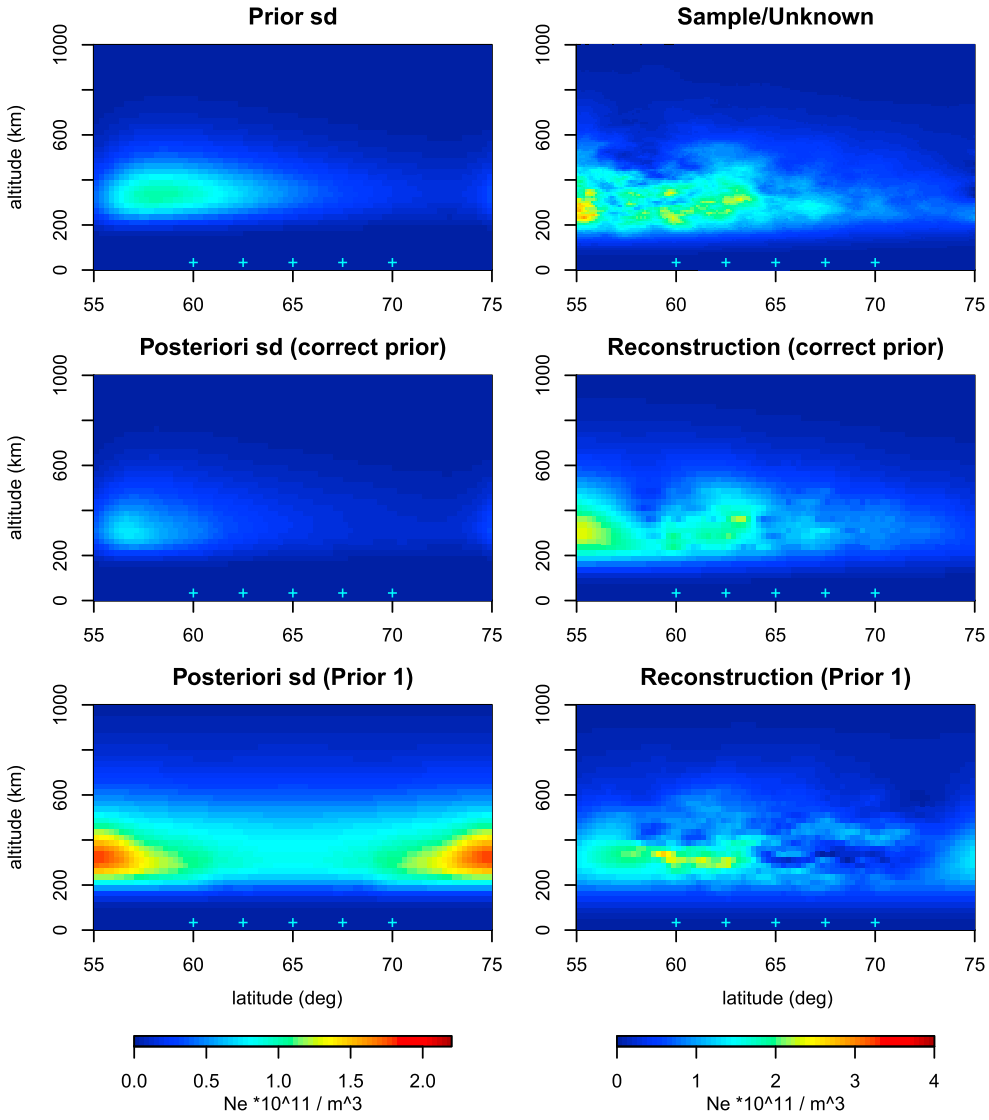

Figure 6. Case 3. (top row) The new prior standard deviation and the corresponding sample. (middle row) The posterior standard deviation and maximum a posteriori estimate with the correct prior. (bottom row) The posterior standard deviation and maximum a posteriori estimate with Prior 1, where the peak value of the mean is decreased to $1.3 \cdot 10^{11}\,\mathrm{Ne/m^3}$. Receivers marked with cyan plus signs.

build the Prior 1 with a large variance so that it would contain only the very basic information we have on the ionosphere.

We then study four different cases illustrated in Figures 4–7. In the first three we form a stricter prior distribution than Prior 1, by dropping the peak standard deviation to half, and generate a random sample from that distribution. The sample is then taken as our unknown electron density. The samples in the example cases are all simulated with the same random seed as sample 1 of Figure 3; hence, the effect of different parameter values can be observed by comparing the samples. In the fourth case in Figure 7, we build a simple electron density with vertical Chapman profile and arbitrary horizontal profile without known prior distribution.

The simulation of measurements is done by solving equation (3) as a forward problem. The domain for the electron density is given in latitudes from 55° to 75°. Vertically, it is limited from ground level to 1000 km altitude. We assume one satellite overflight and five receivers spread uniformly between the latitudes of 60° and 70°. The elevation angle limit for the observations that are used in the tomographic inversion is set to 10°. Given this measurement geometry, we first form the matrix **A**, take the simulated sample as **x**,

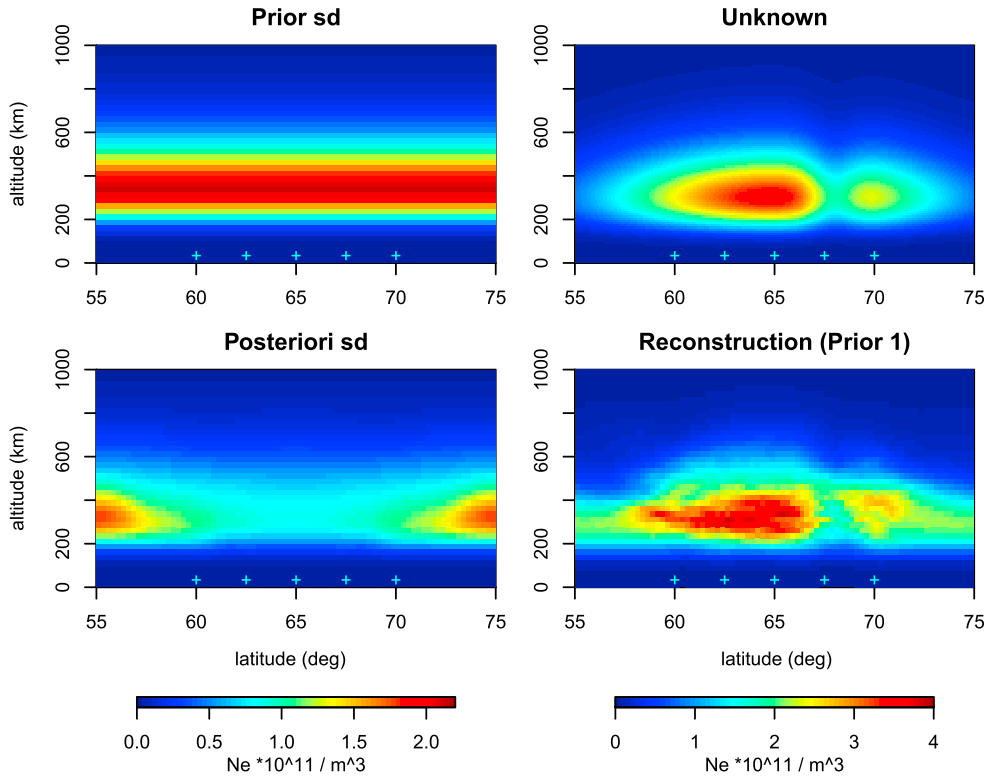

Figure 7. Case 4. (top row) The standard deviation of Prior 1 and a simple synthetical sample. (bottom row) The posterior standard deviation with Prior 1 and the maximum a posteriori estimate. Receivers marked with cyan plus signs.

and compute **m** of equation (2). We then simulate the measurement noise **e** as white noise, with standard deviation $\sigma = 0.01 \max(\mathbf{m})$ and add this to the measurements to get the relative total electron content measurements **M** of equation (3). As explained after equation (1), the real-life ground measurements are relative measurements. Hence, for each satellite-receiver combination we also add an additional random constant γ_i to model the absolute level of total electron content. Here we generate these constants from a Gaussian distribution with a standard deviation of $0.1 \max(\mathbf{m})$. As we are assuming one satellite overflight, this gives us five additional unknown parameters, one for each receiver.

From the simulated measurements, we then solve the "unknown" electron density and phase ambiguity parameters with the estimators given in equations (4) and (5). For the three first cases, the reconstructions are done with both, the correct prior used for sampling and the less restrictive Prior 1 distribution. For the fourth case, the reconstruction is done only with the Prior 1 distribution. In Figures 4–7 the posterior standard deviations are also presented. The posterior standard deviations show the uncertainties of the solution, given the prior distribution and the measurements.

The samples for electron densities are done for 100×200 lattices that correspond to 10 km $\times$ 0.1° pixel size in this example. For the reconstructions we have considered 40×80 lattices with 25 km $\times$ 0.25° pixels.

Case 1. In Figure 4 the scale height of the prior standard deviation is set to 100 km and the peak value to $1 \cdot 10^{11}$ Ne/m^3. The scale height of mean value is also decreased to 125 km. Other parameters are left as in Prior 1. When comparing the sample produced with these parameters to "sample 1" from Figure 3, it can be seen that the main structures remain intact, but the high density is limited to a thinner layer and the variation is lower.

In reconstruction with the correct prior the main structures are recovered, peak densities being slightly lower than in the original electron density. The posterior standard deviation decreases approximately to half from the prior. In reconstruction with Prior 1 the main structures are still somewhat recovered but the overall quality is weaker. As the Prior 1 is less informative, the posterior standard deviation is larger than with the correct prior.

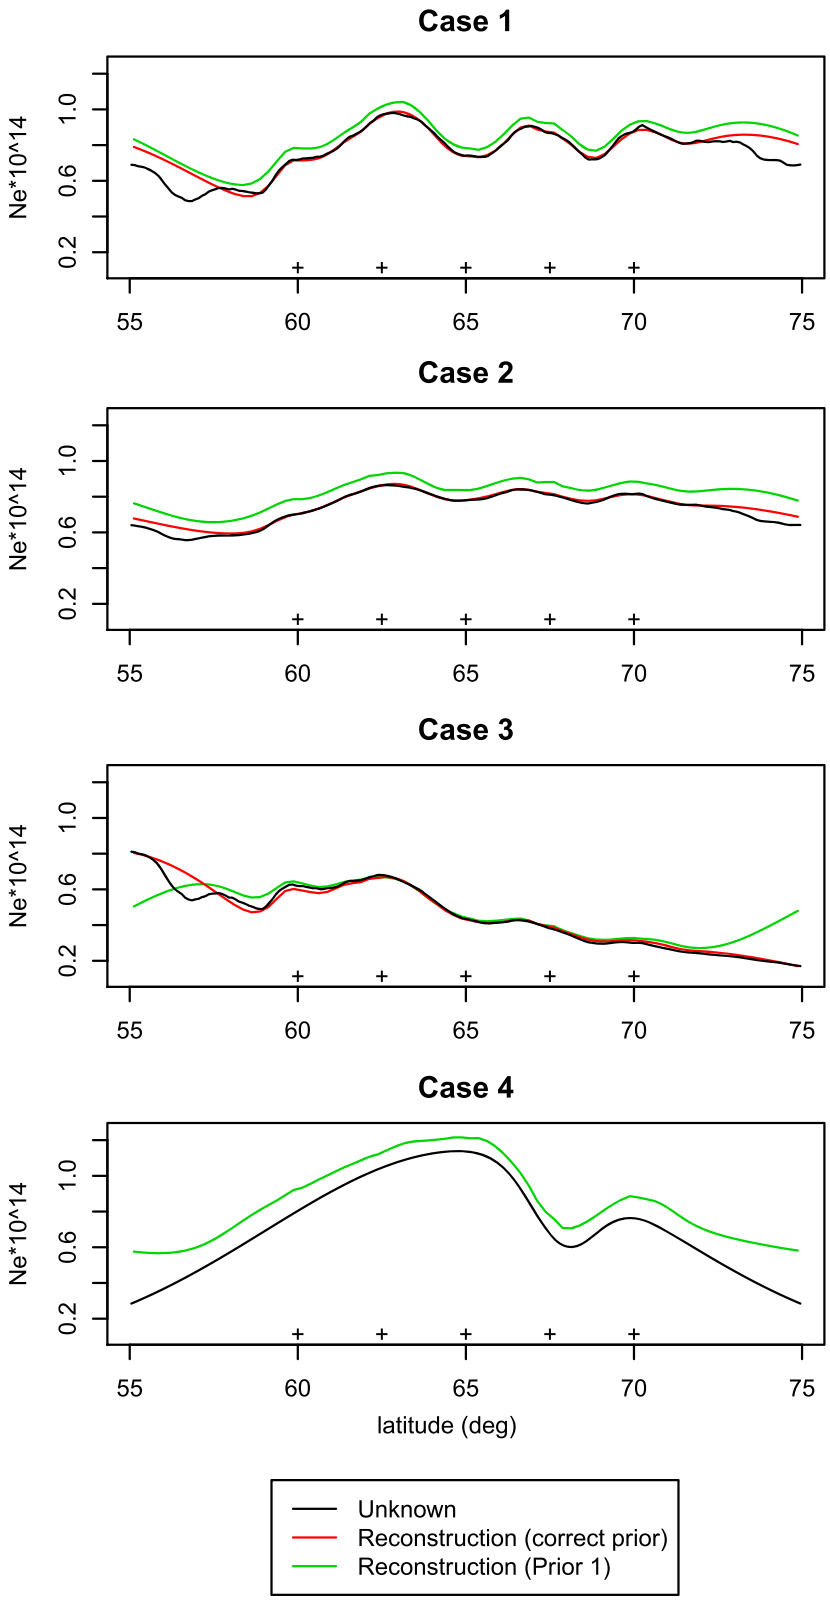

Figure 8. Comparisons of vertically integrated electron content of the original electron density (black) and reconstructions with correct prior (red) and Prior 1 (green). Receivers marked with plus signs.

Case 2. In Figure 5 the parameters are unchanged from Case 1 except for the horizontal correlation length which is set to 20°. When comparing the sample to the sample of the previous case, the effect is clearly visible as the horizontal correlation has increased. With the correct prior, the main features of the reconstructed electron density are reasonably visible, but smoothing is slightly larger than in the previous case. As the horizontal correlation length of Prior 1 is 10°, the Prior 1 can be considered even less informative for Case 2 than it is for Case 1. However, part of the main features is still visible in this reconstruction. As the Prior 1 was used and the measurement geometry was unchanged, the information obtained from the system remains the same as in the previous case, and therefore, the posterior standard deviation is also unchanged.

Case 3. In Figure 6 the parameters are again the same as in Case 1, but the profiles for the prior standard deviation and prior mean have maximum latitude at 57° from where they are horizontally decreasing. The results behave accordingly. As the standard deviation is decreasing northward we are giving stricter information to that direction and the reconstruction with correct prior is therefore also more accurate. This can be seen from the posterior standard deviation which gets very low northward when the correct prior is used. In the first reconstruction trials with Prior 1, in large parts of the area the mean value is significantly larger than in the original electron density; also, some negative densities were produced in reconstruction. Hence, for the reconstruction with Prior 1 in Figure 6 the mean of Prior 1 was scaled down to $1.3 \cdot 10^{11}$ Ne/m^3 until the reconstruction was strictly positive. The posterior standard deviation remains the same compared to previous cases.

Case 4. In Figure 7, a simple electron density is constructed by multiplying a vertical Chapman profile with an arbitrary two-peak horizontal profile. Hence, in this case there is no exact prior knowledge and we rely only on the Prior 1 distribution. In the reconstruction, the main features of the electron density are recovered rather well; however, some values are present that are higher than those in the original electron density. Similar to previous cases, the posterior standard deviation is unchanged as the information provided by the system is the same.

In Figure 8, we have compared the vertically integrated electron content between the original electron density and different reconstructions. In all of the cases the Prior 1 seems to overestimate the electron content slightly. Overall the plots show reasonable agreement between integrated electron contents. In *Vierinen et al.* [2014] the method is applied successfully also to real measurements.

6. Conclusion and Discussion

It is well known that the ionospheric tomography problem, as presented here, is ill posed and cannot be solved uniquely. In order to get a reasonable tomographic reconstruction, additional information is needed to stabilize/regularize the problem. To give this additional information in an easily interpretable manner, we have used Bayesian statistical inversion, where the stabilization is given as a prior distribution.

As the Bayesian approach results in computational difficulties with the full covariance matrices that need to be inverted, we have used GMRF priors to overcome these computational difficulties. We have demonstrated that with GMRFs we can incorporate physically interpretable information in the prior distribution. Most importantly, the GMRFs are implemented with extremely sparse matrices. The sparse matrices provide computational efficiency making the Bayesian statistical inversion approach feasible with simple matrix computations.

In Bayesian statistical inversion, when working with real data, the prior distribution can be chosen based on other measurements or models. However, our uncertainty on this guess should be then reflected in the variance of the prior distribution. The possibility to choose the prior somewhat subjectively might feel unsound, but we stress that all the different tomographic algorithms need the stabilizing/regularizing information and in this Bayesian approach this information is just made visible and interpretable.

In synthetic test cases, where the true electron density is known, the correspondence between the individual reconstruction and the underlying truth can be enhanced considerably by tuning the parameters through trial and error. Instead of using such an approach, the examples presented in section 5 give an overall understanding of how well the ionospheric tomography problem can be solved in extreme cases, with regard to the scale of details in the electron density and to the uncertainty of the state of the ionosphere before the measurements.

In the Bayesian framework, the posterior covariance also provides us with understanding of the uncertainties of the resulting reconstruction. In cases with the correct prior, credible intervals can be computed from the posterior standard deviations. As the solution is highly dependent on the prior in use, one has to be cautious before jumping to conclusions about credible intervals when using real data. However, when prior modeling is done carefully, even this kind of reasoning can be justified.

7. Future Work

We are currently building a new ionospheric tomography receiver network in Fenno-Scandinavia. The future work will include applications of the method presented in this paper to real dual-frequency beacon data. The framework presented in the paper is also well suited for multidata source inversion, allowing the inclusion of measurements from other ionospheric remote sensing instruments, such as ionosondes, GPS receivers, and incoherent scatter radars. We also plan to extend the method for three-dimensional ionospheric tomography.

Acknowledgments
This work has been funded by Academy of Finland (grants 132694 and 250215) and European Regional Development Fund (Regional Council of Lapland, grant A31221). The data used in this study are simulated in R environment and can be regenerated by following the examples.

References

Andreeva, E. S. (1990), Radio tomographic reconstruction of ionisation dip in the plasma near the Earth, *J. Exp. Theor. Phys.*, *54*, 142–148.
Arikan, O., F. Arikan, and C. B. Erol (2007), 3-D Computerized ionospheric tomography with random field priors, in *Mathematical Methods in Engineering*, pp. 325–334, Springer, Netherlands.
Austen, J. R., S. J. Franke, and C. H. Liu (1988), Ionospheric imaging using computerized tomography, *Radio Sci.*, *23*(3), 299–307.
Bardsley, J. M. (2013), Gaussian Markov random field Priors for inverse problems, *Inverse Problem Imaging*, *7*(2), 397–416.
Brekke, A. (1997), *Physics of the Upper Polar Atmosphere*, Wiley, West Sussex.
Bust, G. S., and C. N. Mitchell (2008), History, current state, and future directions of ionospheric imaging, *Rev. Geophys.*, *46*, RG1003, doi:10.1029/2006RG000212.
Davies, K. (1965), *Ionospheric Radio Propagation*, U.S. Dep. of Commer., Washington, D. C.
Fremouw, E. J., J. A. Secan, and B. M. Howe (1992), Application of stochastic inverse theory to ionospheric tomography, *Radio Sci.*, *27*(5), 721–732.
Gentle, J. E. (2005), *Random Number Generation and Monte Carlo Methods—Statistics and Computing*, Springer, New York.
Hansen, A. J. (2002), Tomographic estimation of the ionosphere using terrestrial GPS sensors, PhD dissertation, Stanford Univ., Stanford, Calif.
Kaipio, J., and E. Somersalo (2005), *Statistical and Computational Inverse Problems*, Springer, New York.
Kaipio, J., V. Kolehmainen, M. Vauhkonen, and E. Somersalo (1999), Inverse problems with structural prior information, *Inverse Problems*, *15*, 713–729.
Kunitsyn, V. E., and E. D. Tereshchenko (2003), *Ionospheric Tomography*, Springer, Berlin.
Markkanen, M., M. Lehtinen, T. Nygrén, J. Pirttilä, P. Helenius, E. Vilenius, E. D. Tereshchenko, and B. Z. Khudukon (1995), Bayesian approach to satellite radiotomography with applications in the Scandinavian sector, *Ann. Geophys.*, *13*, 1277–1287.
Mitchell, C. N., and P. S. J. Spencer (2003), A three-dimensional time-dependent algorithm for ionospheric imaging using GPS, *Ann. Geophys.*, *46*, 687–696.
Na, H., and H. Lee (1990), Resolution analysis of tomographic reconstruction of electron density profiles in the ionosphere, *Int. J. Imaging Syst. Technol.*, *2*, 209–218.
Na, H., and E. Sutton (1994), Resolution Analysis of ionospheric tomography systems, *Int. J. Imaging Syst. Technol.*, *5*, 169–173.
Natterer, F. (1986), *The Mathematics of Computerized Tomography*, Wiley, Cichester, U. K., and New York.
Pryse, S. E., L. Kersley, D. L. Rice, C. D. Russell, and I. K. Walker (1993), Tomographic imaging of the ionospheric midlatitude trough, *Ann. Geophys.*, *11*, 144–149.
Raymund, T. D. (1995), Comparisons of several ionospheric tomography algorithms, *Ann. Geophys.*, *13*, 1254–1262.
Raymund, T. D., J. R. Austen, S. J. Franke, C. H. Liu, J. A. Klocbuchar, and J. Stalker (1990), Application of computerised tomography to the investigation of the ionospheric structures, *Radio Sci.*, *25*, 771–789.
Roininen, L., M. Lehtinen, S. Lasanen, M. Orispää, and M. Markkanen (2011), Correlation priors, *Inverse Problem Imaging*, *5*(1), 167–184.
Roininen, L., P. Piiroinen, and M. Lehtinen (2013), Constructing continuous stationary covariances as limits of the second-order stochastic difference equations, *Inverse Problem Imaging*, *7*(2), 611–647.
Roininen, L., J. M. J. Huttunen, and S. Lasanen (2014), Whittle-Matérn priors for Bayesian statistical inversion with applications in electrical impedance tomography, *Inverse Problems Imaging*, *8*(2), 561–586.
Rue, H., and L. Held (2005), *Gaussian Markov Random Fields—Theory and Applications*, Chapman and Hall/CRC, Boca Raton, Fla.
Saksman, E., T. Nygrén, and M. Markkanen (1997), Ionospheric structures invisible in satellite tomography, *Radio Sci.*, *32*(2), 606–616.
Seemala, G. K., M. Yamamoto, A. Saito, and C.-H. Chen (2014), Three-dimensional GPS ionospheric tomography over Japan using constrained least squares, *J. Geophys. Res. Space Physics*, *119*, 3044–3052, doi:10.1002/2013JA019582.
Vierinen, J., J. Norberg, M. S. Lehtinen, O. Amm, L. Roininen, A. Väänänen, P. J. Erickson, and D. McKay-Bukowski (2014), Beacon satellite receiver for ionospheric tomography, *Radio Sci.*, *49*, 1141–1152, doi:10.1002/2014RS005434.
Yeh, K. C., and T. D. Raymund (1991), Limitations of ionospheric imaging by tomography, *Radio Sci.*, *26*(6), 1361–1380.