

Nonlinear regression method for estimating neutral wind and temperature from Fabry–Perot interferometer data

Brian J. Harding,* Thomas W. Gehrels, and Jonathan J. Makela

Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA

*Corresponding author: bhardin2@illinois.edu

Received 16 October 2013; revised 16 December 2013; accepted 20 December 2013;
posted 23 December 2013 (Doc. ID 199601); published 28 January 2014

The Earth's thermosphere plays a critical role in driving electrodynamic processes in the ionosphere and in transferring solar energy to the atmosphere, yet measurements of thermospheric state parameters, such as wind and temperature, are sparse. One of the most popular techniques for measuring these parameters is to use a Fabry–Perot interferometer to monitor the Doppler width and breadth of naturally occurring airglow emissions in the thermosphere. In this work, we present a technique for estimating upper-atmospheric winds and temperatures from images of Fabry–Perot fringes captured by a CCD detector. We estimate instrument parameters from fringe patterns of a frequency-stabilized laser, and we use these parameters to estimate winds and temperatures from airglow fringe patterns. A unique feature of this technique is the model used for the laser and airglow fringe patterns, which fits all fringes simultaneously and attempts to model the effects of optical defects. This technique yields accurate estimates for winds, temperatures, and the associated uncertainties in these parameters, as we show with a Monte Carlo simulation. © 2014 Optical Society of America

OCIS codes: (010.0280) Remote sensing and sensors; (120.2230) Fabry-Perot; (100.2650) Fringe analysis; (280.4991) Passive remote sensing; (300.2140) Emission.

<http://dx.doi.org/10.1364/AO.53.000666>

1. Introduction

The Fabry–Perot interferometer (FPI) is an important tool for upper-atmospheric research, providing remote measurements of neutral wind and temperature. Knowledge of these quantities is important for understanding the dynamics and coupling between the intermingled neutral and plasma environment of the upper atmosphere. Through collisions with ionized species composing the ionosphere, the motion of the neutral species (referred to as the neutral wind) can drive currents and electric fields in the ionosphere [e.g., 1]. The direction and magnitude of the neutral wind is also thought to play a role in the development of instabilities in the ionosphere

[e.g., 2] that can adversely affect technological systems reliant on trans-ionospheric radio wave propagation, such as communication and navigation satellite systems. Furthermore, the response of the neutral atmosphere to geomagnetic activity has an impact on the energy balance of the upper atmosphere through Joule heating effects and driving significant increases in the neutral wind [e.g., 3].

The FPI infers the neutral wind and temperature from the Doppler shift and thermal Doppler broadening of naturally occurring airglow emissions in the upper atmosphere. The specific emission wavelength measured determines the altitude to which the estimate pertains. For measurements of the thermospheric neutral wind, for example, the emission centered at 630.0 nm is used. This emission is caused by the dissociative recombination of O_2^+ and originates from a layer at approximately 250 km in altitude.

The FPI uses a narrowband filter to reduce interference and feeds the light wave through an etalon, generating multiple phase-shifted copies of the wave, which interfere through the lens to create Fabry–Perot fringes on the CCD. In order to determine Doppler shift and breadth from the raw Fabry–Perot fringe pattern, we must account for instrumental effects. Specifically, we must establish the Doppler reference (for wind) and determine the instrumental broadening (for temperature). Often, this calibration step is accomplished by observing a monochromatic source at a known wavelength that fills the field of view, yielding a measurement of the instrument function, which is described in more detail in Section 2.

Various methods have been used to analyze FPI fringes in order to obtain wind and temperature. Perhaps the simplest is to fit Gaussian functions to both the calibration fringes and the airglow fringes [4]. A Gaussian, or indeed any symmetric function, works suitably well for estimating wind because it accurately determines the line center. However, since neither the calibration fringes nor the airglow fringes are actually Gaussians, this technique returns inaccurate temperature. As is well known, the ideal shape of the calibration fringe is an Airy function [5], and using an Airy function instead of a Gaussian has been shown to yield a 120 K difference in the recovered temperature [6], which is an order of magnitude greater than the typical statistical uncertainty in our systems.

A further complication in FPI analysis is that all instruments have defects in the etalon and lens that cause the instrument function to deviate from the ideal Airy function. An elegant solution to this problem is the use of a low-order Fourier fit to the calibration fringes, which can account for such deviation [7]. However, this method treats each fringe independently, which can lead to biases in temperature, presumably due to interference between adjacent fringes, and inaccurate uncertainty estimates [8]. Some researchers use a combination of techniques; for example, some use Gaussian fits to estimate wind and Fourier fits to estimate temperature [9]. For all-sky Doppler imaging systems, an Airy function with optical distortion parameters has been used to account for noncircular fringes [10]. Although these extra distortion parameters are necessary to estimate accurate wind, accurate temperature estimation remains a challenge because these parameters only account for the noncircular nature of the fringe peaks, and not the deviation of the fringe shapes from Airy functions.

As a result of these difficulties in estimating temperature, some researchers attempt to remove this bias in postprocessing by normalizing temperature measurements to an outside source such as an incoherent scatter radar [6]. However, it is our belief that with proper processing, accurate temperature can be obtained directly from the data.

We propose a technique that simultaneously fits all laser fringes as well as a shift-variant

point-spread function. The point-spread function accounts for optical defects in the etalon and lens, leading to accurate estimates of wind and temperature. We then use the inferred instrument function to simultaneously fit all fringes in the sky image. As a byproduct of posing the problem as a model fit, we obtain accurate uncertainties of wind and temperature.

Section 2 describes background material and introduces our models for the calibration fringes and airglow fringes. Section 3 provides details of the inversion procedure, and Section 4 shows results from Monte Carlo simulations, which test this procedure. Section 5 concludes and outlines future work.

2. Background and Forward Model

The heart of the FPI is the etalon, which comprises two parallel, reflective plates of glass separated by an air gap. When illuminated by a source, the etalon creates an interference pattern, which is imaged onto a CCD by an objective lens. As the intent of this paper is for general purpose, we refer the reader to the literature to find more details about the specific instrument used for examples in this paper (MiniME) [8,11,12]. An example of the FPI fringe pattern obtained by an imaging FPI is shown in Fig. 1. We perform an annular summation around the fringe center to reduce this 2D image to a 1D function of radius, r (described in detail in Section 3). Each pixel in this fringe pattern is governed by a Fredholm integral equation of the first kind [8]:

$$S(r) = \int_{-\infty}^{\infty} A(r, \lambda) Y(\lambda) d\lambda, \quad (1)$$

where r is the radial distance from the center of the fringe pattern (usually measured in pixels), λ is a wavelength, $S(r)$ is the pixel count at radius r , $A(r, \lambda)$ is the kernel, known as the instrument function, and $Y(\lambda)$ is the unknown source spectrum (which includes the effects of possible narrowband interference filters used to isolate the atmospheric emission of interest). We have disregarded optical distortions by approximating the fringes as perfectly




Fig. 1. Example FPI airglow fringe image.

circular in order to describe S and A as functions of the radius r . We note that in our systems such optical distortions are not evident, but our model can be extended to include the effects of distortions if necessary, using the techniques developed in [10].

A. Model for Laser Calibration Fringes

In order to estimate spectral information, $Y(\lambda)$, from the raw data, $S(r)$, we need to know the instrument function, $A(r, \lambda)$. In the ideal case, the instrument function is an Airy function [5]:

$$A(r, \lambda) = \frac{I}{1 + \frac{4R}{(1-R)^2} \sin^2\left(\frac{2\pi n t}{\lambda} \cos \theta(r)\right)}, \quad (2)$$

where I is called the intensity and has units of CCD counts, R is the reflectivity of the etalon plates, n is the index of refraction in the etalon gap, t is the etalon gap in meters, and $\theta(r)$ is the angle with the optical axis. Pixels at larger radial distances on the CCD capture light from larger angles. For the FPI used as an example in this paper, the field of view is 1.5° . For narrow-field FPIs such as this one, the function $\theta(r)$ is given by

$$\theta(r) = \tan^{-1}(\alpha r), \quad (3)$$

where α is a constant that will be estimated from the laser calibration image. For use in wide-angle FPI systems, a more suitable representation of the lens may need to be substituted.

In practice, all etalons and lenses will have defects, causing the instrument function to deviate from this ideal Airy function by broadening each fringe. Early work resulted in analytical broadening functions for a number of defects including microscopic etalon flatness imperfections, spherical defect of the etalon plates, and the finite aperture [13]. However, these analytical expressions are not enough to fully characterize the instrument function, so a calibration using a monochromatic source is necessary [7]. We see from Eq. (1) that a monochromatic source spectrum [i.e., a delta function for $Y(\lambda)$] will give us an




Fig. 2. Example FPI laser calibration fringe image.

image that is a copy of the instrument function at that wavelength. For the monochromatic source, we use a 632.8 nm frequency-stabilized HeNe laser, since it is reasonably close in wavelength to the 630.0 nm airglow emission we observe.

We show an example laser calibration image in Fig. 2. We perform an annular summation on this raw image to obtain a 1D function of radius, r , which is described in more detail in Section 3. This reduced fringe pattern is shown in Fig. 3. The data are plotted as a function of r^2 so that the fringes appear equally spaced. Plotted with the data is the best-fit ideal Airy function. We also show the residual, the difference between the data and the Airy model. At least two deviations from the ideal Airy function are clear. First, vignetting, absorption, and other nonidealities cause attenuation of the fringes near the edge of the CCD. Second, the shape of the fringes appears to be shorter and wider than an ideal Airy function, presumably due to the broadening described above.

In order to better characterize the instrument function, we introduce additional parameters to our model to account for these defects. First, we replace the constant I in Eq. (2) with a quadratic falloff term:

$$I = I_0 \left(1 + I_1 \left(\frac{r}{r_{\max}} \right) + I_2 \left(\frac{r}{r_{\max}} \right)^2 \right), \quad (4)$$

where r_{\max} is the maximum value of r on the CCD. In practice, I_1 and I_2 are often negative. We have tried higher-order fits but find that a quadratic function is sufficient. This accounts for the first deviation described above.

The second deviation is accounted for by introducing into our forward model a point-spread function, $b(s, r)$. This point-spread function blurs the Airy pattern to create a modified Airy function, $\tilde{A}(r)$:

$$\tilde{A}(r, \lambda) = \int_0^{r_{\max}} b(s, r) A(s, \lambda) ds, \quad (5)$$

where the integral is taken over the entire fringe pattern, and s is a dummy variable in the same coordinates as r . Each pixel of the modified Airy function, $\tilde{A}(r, \lambda)$, is a weighted average of the neighborhood




Fig. 3. Example instrument function with best-fit ideal Airy function. The difference (residual) between data and model is also shown. Optical defects cause a large mismatch, which demands a more complex model.

pixels of the ideal Airy function, where the weights are determined by the shape of the point-spread function at that pixel, $b(s, r)$. The effect of this blurring is to broaden the fringes, which simulates the effects of optical defects.

In a first approximation, we assume the width of the point-spread function does not change across the CCD and takes the form of a Gaussian centered on the pixel, r :

$$b(s, r) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(s-r)^2}{\sigma^2}}. \quad (6)$$

This point-spread function is normalized so that its integral is unity, and thus its width, σ , is its only free parameter. Note that such a constant-width Gaussian will blur the outer fringes more than the inner fringes since the outer fringes are closer together, which matches what we see in practice.

However, because of varying path lengths through the optics, the point-spread function, $b(s, r)$, is not constant across the CCD, and its width, σ , varies slightly. As such, we have found it useful to allow σ to vary as a function of r , so we update Eq. (6):

$$b(s, r) = \frac{1}{\sqrt{2\pi\sigma(r)^2}} e^{-\frac{(s-r)^2}{\sigma(r)^2}}, \quad (7)$$

where we parameterize $\sigma(r)$ with a first-order Fourier series:

$$\sigma(r) = \sigma_0 + \sigma_1 \sin\left(\pi \frac{r}{r_{\max}}\right) + \sigma_2 \cos\left(\pi \frac{r}{r_{\max}}\right). \quad (8)$$

Unlike our fit of intensity, I , we chose to implement a Fourier series fit instead of a polynomial fit for a numerical reason: different Fourier series terms are orthogonal, while polynomial terms are not. Choosing a Fourier series fit reduces correlation between σ_0 , σ_1 , and σ_2 , allowing us to achieve convergence more reliably. That being said, we find values of σ_1 and σ_2 often smaller than σ_0 , so other parameterizations of $\sigma(r)$ are likely possible.

Table 1. Parameters of the Forward Model for Laser Calibration Fringes*

Parameters	Description	Fixed/ Varied	Typical Value
λ	Laser wavelength	Fixed	632.8 nm
n	Index of refraction	Fixed	1
R	Reflectivity	Varied	0.8
t	Etalon gap	Varied	15 mm
α	Magnification constant	Varied	8.5×10^{-5}
I_0	Average intensity	Varied	1000 counts
I_1	Linear falloff of intensity	Varied	-0.1
I_2	Quadratic falloff of intensity	Varied	0.005
σ_0	Average blur size	Varied	0.8 pixels
σ_1	Sin-variation of blur size	Varied	0.1 pixels
σ_2	Cos-variation of blur size	Varied	-0.05 pixels
B	CCD bias	Varied	300 counts

*Some parameters are fixed and some are allowed to vary to fit the data.

We now update our forward model, Eq. (1), with the modified Airy function and add the CCD bias, B , as a parameter:

$$S(r) = \int_{-\infty}^{\infty} \tilde{A}(r, \lambda) Y(\lambda) d\lambda + B. \quad (9)$$

Taken together, Eqs. (2)–(5) and (7)–(9) define a forward model for the calibration laser fringe pattern. The parameters in this model and their typical values for our systems are shown in Table 1. More details on how these parameters are estimated from the raw data are described in Section 3. An example model fit is shown in Fig. 4, where it can be seen that the residual is much smaller than in Fig. 3. However, the fit is not perfect, with a χ^2 value of 2.6. This is due to two facts: our model does not perfectly capture the effects of optical distortions, and the high signal-to-noise ratio causes small deviations from the model to significantly impact χ^2 .

B. Model for Airglow Fringes

With the model for the laser calibration fringes developed, we turn to the model for the airglow fringes. These are of course governed by the same equations, except the source spectrum, $Y(\lambda)$, is no longer a delta function but is assumed to be a thermally broadened and Doppler-shifted Gaussian, the equation for which is given in [8] and is reproduced here with a typographical error corrected:

$$Y(\lambda) = Y_{\text{bg}} + Y_{\text{line}} \exp\left\{-\frac{1}{2}\left(\frac{\lambda - \lambda_c}{\Delta\lambda}\right)^2\right\}, \quad (10)$$

where Y_{bg} is the background sky emission, assumed to be independent of λ , Y_{line} is the intensity of the airglow line, λ_c is its center wavelength, and $\Delta\lambda$ is its Doppler breadth. The line center is related to the line-of-sight wind velocity through a Doppler shift:

$$\lambda_c = \lambda_0 \left(1 + \frac{v}{c}\right), \quad (11)$$

where λ_0 is the nominal line center, v is the line-of-sight velocity away from the instrument, and c is




Fig. 4. Example instrument function with best-fit modified Airy function, and residual between data and modified Airy fit. Compare to Fig. 3.

Table 2. Free Parameters of the Forward Model for Airglow Fringes^a

Parameters	Description
B	CCD bias
Y_{bg}	Background emission
Y_{line}	Line intensity
λ_c	Line center
$\Delta\lambda$	Linewidth

^aAll other parameters are fixed at the values determined by the laser calibration analysis.

the speed of light. The Doppler broadening is related to the temperature through [8]:

$$\Delta\lambda = \frac{\lambda_0}{c} \sqrt{\frac{kT}{m}}, \quad (12)$$

where k is the Boltzmann constant, T is the temperature of the emitting species, and m is its mass.

Taken together, Eqs. (2), (3), (4), (5), (7), (8), (9), and (10) define a forward model for the airglow fringes. The free parameters are shown in Table 2 and are varied to match the airglow fringes, a process that is described in more detail in Section 3. All other parameters are fixed at the values found by analyzing the laser calibration fringes. An example model fit of an airglow fringe pattern is shown in Fig. 5.

3. Inversion

In this section, we describe the details of the data reduction and inversion. First, we reduce the 2D calibration and airglow fringe images to 1D fringe patterns. Next, we fit the model described in Section 2.A to the calibration fringes in order to estimate instrument parameters, and we use these instrument parameters to fit the model in Section 2.B to the airglow fringes.

Although it is possible to fit our models directly to the 2D fringe images, we first reduce our data to 1D fringe patterns to decrease computation time. The details of this reduction are provided in the literature [8] and are briefly outlined here. First, we determine the fringe center from the laser calibration image by thresholding the image, identifying contiguous regions as the fringes, fitting circles to each fringe,




Fig. 5. Example airglow fringes with best-fit airglow model. Also shown is the residual, the difference between the data and the model.

and taking the median center value. This yields an estimate of the center pixel with a typical statistical variation of a few hundredths of a pixel, which is less than the 0.3 pixels required to maintain a temperature error below 1% [14]. With the center determined, the pixels are sorted as a function of radius and averaged into R radial bins, a process known as annular integration. The bin locations are chosen so that all bins contain an approximately equal number of pixels. For our 512×512 -pixel images, we use $R = 500$. Our simulations for our systems suggest that as long as the number of bins is greater than about 200 yet small enough that each bin contains several pixels, the inversion is not sensitive to the number of bins.

We denote the average pixel count in the radial bin with radius r as $N_{\text{laser}}(r)$ and $N_{\text{airglow}}(r)$, for the laser calibration and airglow fringes, respectively. We proceed with the laser calibration fit by the standard procedure of choosing the instrument parameters that minimize χ^2 :

$$\chi^2 = \sum_{i=0}^{R-1} \left(\frac{N_{\text{laser}}(r_i) - S(r_i)}{\sigma_{N_{\text{laser}}}(r_i)} \right)^2, \quad (13)$$

where r_i is the radius of the i th radial bin, there are $R = 500$ radial bins, and $S(r_i)$ is from Eq. (9), the forward model for the calibration fringes, described in Section 2.A. The statistical uncertainty of $N_{\text{laser}}(r_i)$, given by $\sigma_{N_{\text{laser}}}(r_i)$, is a bootstrapped estimate, found by dividing the standard deviation of the pixel counts in a radial bin by the square root of the number of pixels in the bin.

In order to numerically evaluate the forward model, $S(r_i)$, we must discretize the two integral equations, Eqs. (5) and (9). For Eq. (5), we do this in the straightforward way of replacing the integral with a summation over the radial bins:

$$\tilde{A}(r, \lambda) = \sum_{i=0}^{R-1} b(s_i, r) A(s_i, \lambda) \Delta s_i, \quad (14)$$

where in practice, b is narrow enough that the Δs_i term is taken as constant. For Eq. (9), a discrete form is not necessary for the laser calibration fringes, since we can analytically integrate over a delta function. However, we will need Eq. (9) for the airglow fringe analysis, so we approximate the integral as a sum over a suitably large wavelength range (typically, five free spectral ranges):

$$S(r_i) = \sum_{j=0}^{L-1} \tilde{A}(r_i, \lambda_j) Y(\lambda_j) \Delta\lambda + B, \quad (15)$$

where we have discretized the spectrum into L equally spaced wavelength bins, λ_j , and we have R copies of Eq. (15), one for each radial bin, r_i . Stacked vertically, these R equations can be written in matrix form as:

$$\mathbf{s} = \mathbf{A}\mathbf{y} + B\mathbf{1}, \quad (16)$$

where \mathbf{s} is a vector of length R containing the model fringe pattern, \mathbf{A} is a $R \times L$ matrix, which has columns containing monochromatic fringe patterns, \mathbf{y} is a vector of length L containing the discretized spectrum, and $\mathbf{1}$ is a length R vector with every element equal to 1. This form offers an intuitive interpretation of the Fredholm integral, Eq. (9): the output of the FPI (\mathbf{s}) is a superposition of monochromatic fringe patterns (columns of \mathbf{A}), weighted by the source spectrum (\mathbf{y}).

To find the instrument parameters that minimize χ^2 , we use the well-known Levenberg–Marquardt algorithm, which is provided in the **lmfit** Python library [15]. This algorithm solves for optimized parameter estimates along with statistical uncertainties in these parameters. However, because the Airy function is highly nonlinear in many of its parameters, care must first be taken to provide an adequate initial guess for these parameters. Some of these parameters, such as I_0 , I_1 , I_2 , and B , can be approximated from the data. Others, such as α and R , are approximately known from the instrument design. For the etalon gap, t , we take advantage of the fact that to the first order, the Airy function is periodic in t , and we know the approximate value of t from the instrument design. We perform a brute-force search over the period, $\lambda/2$, to find an adequate initial guess for t . For the remaining parameters, σ_0 , σ_1 , and σ_2 , we find that small initial guesses ($\sigma_0 = 0.5$ pixels, $\sigma_1 = \sigma_2 = 0$ pixels) are sufficient.

Instead of immediately optimizing all the parameters simultaneously (a strategy that can lead to a non-optimal local minimum), we proceed in stages, optimizing only certain parameters at a time in order to obtain a better initial guess for the next stage. The final step is a full optimization over all parameters. The ad-hoc sequence of stages was chosen through testing and in our simulations appears to ensure convergence. This approach is overly conservative because our data set is too large to allow for human oversight, so the inversion must succeed every time.

With the instrument parameters determined, we move on to analyze the airglow fringes. The same process described above is applied. We use the center location found from the laser calibration to average the data into radial bins. We then perform a fit of the model described in Section 2.B, discretized as described above. For the discretization in Eq. (15), we use $L = 101$, which our simulations show is large enough to achieve accurate airglow parameters. The fit minimizes a χ^2 function analogous to Eq. (13). Again, we must take care to find good initial guesses for the parameters. Two parameters can be directly estimated from the data: Y_{line} and B . Much like for the etalon gap, the initial guess for the line center, λ_c , is estimated with a brute-force search over one free spectral range. The line width, $\Delta\lambda$, is initially set so that $T = 1000$ K, and the background emission, Y_{bg} , is initially set to zero. As above, we use

the Levenberg–Marquardt algorithm and optimize in stages in order to ensure convergence. Finally, we convert the airglow parameters λ_c and $\Delta\lambda$ to line-of-sight neutral velocity and temperature (and the associated uncertainties) using Eqs. (11) and (12).

A practical detail in these experiments is that the instrument cannot measure airglow and calibration fringes simultaneously. In order to allow the instrument parameters to vary over time, the instrument parameters used to invert an airglow image must be interpolated in time from laser calibration exposures before and after the airglow exposure.

There are three other practical details of possible concern. First, the laser calibration images are corrupted by noise, yielding noisy instrument parameter estimates, yet these parameters are assumed exact in the airglow image inversion. However, since the laser is much brighter than airglow, this effect is small, as shown in Section 4.A.

Second, we incorrectly assume that the laser wavelength is known exactly, which leads to a constant offset in our estimated line-of-sight wind. We remove this offset by assuming the average vertical wind over the night is zero, which is a less restrictive assumption than the historical practice of assuming its instantaneous value is zero [8].

The final concern relates to the estimate of the etalon gap, t , from the laser fringes. In the Airy function, this parameter is present only in the term nt/λ , so we must assume that the index of refraction, n , and the laser wavelength, λ , are constant. Temporal variations in the index of refraction or the laser wavelength map directly into variations in the estimated wind. In our systems, the laser frequency stability is ± 1 MHz over 8 h, which translates to a velocity drift of ± 0.6 (m/s), and we do not expect appreciable drifts in the index of refraction. Any constant offset in n or λ will be corrected for by assuming the average vertical wind is zero.

4. Monte Carlo Simulations

In order to test this inversion method, we perform three Monte Carlo simulations. The first simulation tests the uncertainty estimates, the second tests for biases over the range of expected velocities and temperatures, and the last tests for biases over signal-to-noise ratio (SNR).

Each Monte Carlo trial consists of simulating truth images for a given “true” wind and temperature, adding realistic noise, and running the inversion routine to recover estimates of wind and temperature. We simulate two truth images: the laser calibration image and the airglow image, using the forward models described in Sections 2.A and 2.B with the discretization given in Section 3. To generate the laser calibration image, we use instrument parameters that match those we typically see in practice (e.g., Table 1). For the airglow image, we vary the true wind and temperature and use values of B , Y_{bg} , and Y_{line} to match what we see in practice. We use a fine spectral

resolution ($L = 300$) to generate the airglow image, but use $L = 101$ in the inversion.

We add Poisson noise to the laser calibration image because this observation is photon-noise limited in practice. We add Gaussian white noise of a given SNR to the airglow image because dark noise tends to dominate. If we define ΔS as the difference between the maximum and minimum CCD counts of $S(r)$, then we define the SNR as

$$\text{SNR} = \frac{\Delta S}{\sigma_N}, \quad (17)$$

where σ_N is the standard deviation of the Gaussian white noise, measured in CCD counts. Finally, we apply the inversion method described in Section 3: estimate the center pixel from the laser image, reduce both images to 1D fringes, and fit our models to obtain estimates of wind and temperature.

A. Simulation 1: Uncertainty Estimates

For our first Monte Carlo simulation, we use constant values of $v = 100$ (m/s), $T = 800$ K, and $\text{SNR} = 5$. We run 10^4 trials and compare our estimated uncertainties with the actual error in the estimated values. The results of all 10^4 trials are plotted in Fig. 6. The estimated values follow the expected Gaussian distribution, with 68% of the dots lying within the error bounds. The blue ellipse represents the sample covariance matrix of v and T , from which we see that there is no discernable correlation between errors in v and T . We also find that typical uncertainties in these parameters for $\text{SNR} = 5$ are $\sigma_v = 1.8$ (m/s) and $\sigma_T = 6.5$ K. The red ellipse represents the average estimated uncertainties. The alignment of the red and blue ellipses indicates that the inversion routine is accurately estimating the uncertainties. This implies that the small amount of noise in the laser calibration images has a negligible effect on the airglow parameter uncertainties. Having verified the statistical errors, we investigate systematic errors in the second and third simulations.




Fig. 6. Recovered wind and temperature for 10^4 simulated FPI measurements. Also shown are the sample and estimated uncertainty ellipses, which compare favorably, indicating accurate estimation of error bars.

B. Simulation 2: Biases Over Wind and Temperature

For our next simulation, we use randomized velocities and temperatures as known inputs into the forward model and a constant $\text{SNR} = 5$. Velocities and temperatures are chosen uniformly at random from -300 to 300 (m/s) and from 300 to 1500 K, respectively. We run 10^3 trials and show the results in Fig. 7.

For each trial we calculate the velocity and temperature error (the difference between estimated value and actual value) and plot it against the true velocity and temperature. For the velocity estimates, we find that for all values of v and T there is a constant 0.4 (m/s) bias of unknown origin. This bias is negligible compared to the bias arising from the unknown laser wavelength, and it is removed by assuming there is no average vertical wind over the night, as discussed in Section 3. For the temperature estimates, we find no statistically significant bias, although we note that uncertainties are higher when temperatures are higher, as expected [8].

C. Simulation 3: Biases Over SNR

Finally, we run a simulation with randomized SNR between 0.5 and 5 and randomized velocities and temperatures as in the previous simulation. We run 10^4 trials and show the results in Fig. 8. For all values of SNR, we find no bias in velocity or temperature.

These three simulations suggest that our inversion method produces no significant biases in velocity and temperature over ranges of velocity, temperature, and SNR expected for real-world data, and that the uncertainty in these parameters is accurately estimated. However, these simulations were performed under the ideal case in which our instrument function is well modeled by our forward models described above. Our method could produce biases for instruments that deviate from these models.




Fig. 7. Estimation errors in velocity and temperature as a function of true velocity and temperature for 10^3 simulated FPI measurements. A small, 0.4 (m/s) velocity bias is present, but there is no bias in temperature. In 68% of the trials, the error bar contains zero (not shown), which matches the expectation from Gaussian statistics.




Fig. 8. Estimation errors in velocity and temperature for 10^4 simulated FPI measurements, as a function of SNR. No biases are evident. In 68% of the trials, the error bar contains zero (not shown), which matches the expectation from Gaussian statistics.

5. Conclusion

We have introduced a technique for estimating wind and temperature from laser calibration and airglow fringe images obtained from a Fabry–Perot interferometer. This technique fits a single model to the entire calibration fringe pattern and includes the effects of optical defects in the instrument function. Our Monte Carlo simulations show that this technique produces no significant biases in estimates of wind or temperature, and the associated uncertainties are accurately estimated.

In the future, we plan to empirically validate this method by comparing observations from different instruments and cross-validating with results obtained from incoherent scatter radars. Although these radars provide estimates of the ion and electron temperatures, under the thermodynamic equilibrium conditions expected in the upper atmosphere at nighttime, they can be directly compared to the neutral temperature estimated from an FPI observing the thermospheric 630.0 nm emission layer. Initial work has indicated that there may be a 30–50 K temperature bias between instruments, which deserves further investigation. We have currently only applied this technique to the MiniME family of FPIs, and we hope to verify it on other instruments in the future.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under grant no. DGE-1144245. Work at

the University of Illinois was also supported by the National Science Foundation CEDAR program under grant no. AGS-1138998. The authors acknowledge useful conversations with Prof. J. W. Meriwether at Clemson University.

References

1. R. Heelis, "Electrodynamics in the low and middle latitude ionosphere: a tutorial," *J. Atmos. Sol. Terr. Phys.* **66**, 825–838 (2004).
2. E. Kudeki, A. Akgiray, M. Milla, J. Chau, and D. L. Hysell, "Equatorial spread-F initiation: Post-sunset vortex, thermospheric winds, gravity waves," *J. Atmos. Sol. Terr. Phys.* **69**, 2416–2427 (2007).
3. T. J. Fuller-Rowell, M. V. Codrescu, R. J. Moffett, and S. Quegan, "Response of the thermosphere and ionosphere to geomagnetic storms," *J. Geophys. Res.* **99**, 3893 (1994).
4. M. Biondi, D. P. Sipler, M. E. Zipf, and J. L. Baumgardner, "All-sky Doppler interferometer for thermospheric dynamics studies," *Appl. Opt.* **34**, 1646–1654 (1995).
5. G. Hernandez, *Fabry-Perot Interferometers*, Cambridge Studies in Modern Optics (Cambridge University, 1988).
6. C. G. M. Brum, C. A. Tepley, J. T. Fentzke, E. Robles, P. T. dos Santos, and S. A. Gonzalez, "Long-term changes in the thermospheric neutral winds over Arecibo: climatology based on over three decades of Fabry-Perot observations," *J. Geophys. Res.* **117**, A00H14 (2012).
7. T. Killeen and P. Hays, "Doppler line profile analysis for a multichannel Fabry-Perot interferometer," *Appl. Opt.* **23**, 612–620 (1984).
8. J. J. Makela, J. W. Meriwether, Y. Huang, and P. J. Sherwood, "Simulation and analysis of a multi-order imaging Fabry-Perot interferometer for the study of thermospheric winds and temperatures," *Appl. Opt.* **50**, 4403–4416 (2011).
9. K. Shiokawa, T. Kadota, M. K. Ejiri, Y. Otsuka, Y. Katoh, M. Satoh, and T. Ogawa, "Three-channel imaging fabry-perot interferometer for measurement of mid-latitude airglow," *Appl. Opt.* **40**, 4286–4296 (2001).
10. M. Conde, "Deriving wavelength spectra from fringe images from a fixed-gap single-etalon Fabry-Perot spectrometer," *Appl. Opt.* **41**, 2672–2678 (2002).
11. K. Shiokawa, Y. Otsuka, S. Oyama, S. Nozawa, M. Satoh, Y. Katoh, Y. Hamaguchi, Y. Yamamoto, and J. Meriwether, "Development of low-cost sky-scanning Fabry-Perot interferometers for airglow and auroral studies," *Earth Planets Space* **64**, 1033–1046 (2012).
12. J. J. Makela, J. W. Meriwether, J. Lima, E. S. Miller, and S. Armstrong, "The remote equatorial nighttime observatory of ionospheric regions project and the international heliospherical year," *Earth Moon Planets* **104**, 211–226 (2009).
13. P. B. Hays and R. G. Roble, "A technique for recovering Doppler line profiles from Fabry-Perot interferometer fringes of very low intensity," *Appl. Opt.* **10**, 193–200 (1971).
14. S. Armstrong, "Fabry-Perot data analysis and simulation for the renoir observatories," Master's thesis, (University of Illinois, 2008).
15. M. Newville, "Non-linear least-square minimization for python," <http://newville.github.io/lmfit-py/> (2013).