AO09: Interpreting Satellite Measurements of Atmospheric Ammonia

Candidate number: 1045056
Supervisor: Dr. A. Dudhia
Word count: 3626

Abstract

The vast majority of atmospheric ammonia (NH_3) originates from agricultural practices. Tracking the production of NH_3 is key to maintaining high air quality and better modelling of climate change. Satellite-based infrared spectroscopy has facilitated the detection of NH_3 sources on a global scale. However, quantifying the magnitude of any given source remains an inaccurate process, owing to NH_3's low atmospheric concentration and thus weak absorption features. This project seeks to tackle these difficulties by providing a new means of interpreting NH_3 retrievals.

The *a priori* assumption of NH_3 profile is highlighted as a key cause of retrieval inaccuracy. Previous fast-retrieval methods are presented, and kernel weighting functions(KWFs) are introduced as an alternative means approaching this inaccuracy. Full derivations are given for both the linear, radiative transfer model and the expected functional form of the KWFs. KWFs are shown to produce retrievals accurate to within 3%, independent of assumed profile. KWFs are tested with changes to thermal contrast and atmospheric temperature profile, with the measured response behaving in accordance with linear, radiative transfer theory. Finally, retrieval uncertainty is calculated for five satellite instruments, capable of NH_3 detection. Best-case estimates suggest a factor of 4 improvement in retrieval uncertainty can be expected from the next generation instrument, IASI-NG, set to launch in late 2024.

1 Introduction

1.1 Motivation and Background

Ammonia (NH_3) is the most prevalent alkali in the atmosphere by volume[1]. Agriculture is largely responsible for NH_3 emissions[2]; specifically, losses from fertiliser treatment and volatilisation of livestock manure are estimated to account for 70% of global emissions. Other anthropogenic NH_3 emitters include fossil fuel combustion and direct industrial emission [3].

Emitted NH_3 readily reacts with acidic pollutants, e.g. sulfuric acid and nitric acid, and the resulting neutralisation product, ammonium (NH_4^+), poses risks to human health[2] as a fine particulate. Given the high degree of regulatory attention already faced by acidic pollutants, targeting NH_3 emission is thought of as a cost-effective approach to reducing the concentration of $PM_{2.5}$ (fine particulate matter)[1]: recent studies estimate a 50% reduction in NH_3 emissions in North-Western Europe as sufficient to reduce $PM_{2.5}$ by 25%.

Ammonia pollution also has environmental implications. Proximal to the emission site, NH_3 can be assimilated back into aquatic and terrestrial ecosystems. Its basic chemistry prompts severe eutrophication, biodiversity loss and excessive nitrogen uptake, diminishing crop yield[4]. Further, NH_4^+ scatters incoming solar radiation, while also acting as a site for cloud nucleation. Combined, these lead to net negative radiative forcing[5], and thus atmospheric NH_3 pollution must be accounted for as part of climate change estimation models.

1.2 Remote Sensing of NH_3

In order to control NH_3 production, spatial and temporal trends in NH_3 emission must be accurately characterised. There are two primary means of quantifying the presence of atmospheric pollutants: *bottom-up* and *top-down* approaches. Bottom-up strategies rely on cataloguing all significant emission sources, and totaling their respective contributions. Top-down methods indirectly measure the presence of NH_3

via physical measurements from space.

Bottom-up approaches remain limited in scope, primarily due to the largely unregulated use of nitrate-based fertilisers and hence inaccurate estimates of surface emissions[6]. In-situ (surface and low-altitude) measurements can help improve such estimates, but are still heavily constrained by NH_3's short lifetime. Accordingly, very few in-situ measurements record with temporal resolutions less than 1 month and instead focus on observing longer-term trends [7].

Top-down analyses have long been employed for highly-abundant, atmospheric gases (carbon dioxide (CO_2), water (H_2O) etc.) which only require low-resolution, broad-band filters. In contrast, it was only after the advent of high-spectral resolution spectrometers in the early 2000s[2], that the relatively weak absorption signal of NH_3 could be detected from space. Since then, top-down measurements have shown greater consistency than bottom-up studies, when mapping regional variations in NH_3[6].

This report will discuss five operational, and planned, Fourier-transform infrared spectroscopy (FTIR) instruments. The most comprehensive NH_3 catalogue has been assembled by the Infrared Atmospheric Sounding Interferometer (IASI)[8]. Secondary data is also available from the Cross-track Infrared Sounder (CrIS)[9] and the Atmospheric InfraRed Sounder (AIRS)[10]. Expanded NH_3 catalogues are expected from next-generation FTIR instruments, IASI-NG[11] and IRS [12]. The relative performance of each instrument is evaluated in section 4.

1.3 Instrumentation and Retrieval Products

To-date, IASI remains our primary instrument for mapping atmospheric NH_3; a more detailed instrument specification is outlined below. Comparable missions, both operational and planned, are also presented in table 1.

This project has three main aims. Firstly, to demonstrate the sensitivity of retrieval accuracy to the assumed vertical ammonia profile. Secondly, to introduce kernel weighting functions – henceforth denoted KWFs – as a method of overcoming the retrieval dependence on the assumed profile. Finally, to quantify the retrieval errors for various satellite instruments, accounting for differences in spectral resolution[1] and instrument noise, as described in table 1.

1.3.1 IASI Specification

IASI operates with a relatively poor spatial resolution of 12 km per pixel, when compared with more moderate spectral imagers (e.g. MODIS[14] ≈ 500 m per pixel). The limited spatial resolution makes the detection of individual NH_3 sources difficult, and is systematic of all high resolution FTIRs. To achieve the high spectral resolution required for NH_3 detection, constraints are imposed on the solid angle (and thus spatial resolution) that can be observed over for a given radiance (see section 2.2.1).

IASI follows a sun-synchronous orbit; the equatorial line is crossed at mean local solar times of 09:30 LT and 21:30 LT, providing full Earth coverage, twice a day. Infrared sounding is temporally most effective when the sun has adequately warmed the Earth's surface, but insufficient time has elapsed to allow this heat to be distributed throughout the atmosphere via convection. Hence, the morning node of observation is considered preferable for atmospheric retrievals.

IASI has generated over 15 years worth of retrieval product, IASI-NH3, contributing to an inventory of observable NH_3 sources[8]. Generating agreement, however, between in-situ, satellite and bottom-up approximations of NH_3 concentration has proved difficult[15]. Some reports assert discrepancies ranging from 20% to 50%[16]. It is hoped that future satellites, e.g. IASI-NG, will be able to improve this disparity.

1.4 Characterising Ammonia Profiles

Each IASI measurement has an associated Hyperspectral Range Index (HRI), characterising the observed radiance spectra. Converting the HRI into a NH_3 column amount is non-trivial, and relies on iteratively performing a full radiative transfer calculation. Two methods have emerged to expedite this process by linking variations in the Jacobian spectra with atmospheric temperature structure: Look-up Ta-

---

[1]AIRS is a grating spectrometer, as opposed to a FTIR spectrometer common to the other instruments. Hence, AIRS measures in uniform increments of wavelength, rather than wavenumber. For purposes of comparison, we state the resolution in terms of wavenumbers, assuming a constant $\tilde{\nu} = 900$ cm^{-1}

Table 1: Infrared Spectroscopy Satellite Instruments

Mission	Satellite	Launch Date	Spectral Resolution (cm^{-1})	NEΔT(K)
AIRS[10]	Aqua Satellte	04/05/02	0.75	0.1
IASI[13]	MetOp-A	19/10/06	0.5	0.15
CrIS[9]	Suomi NPP	28/10/11	0.625	0.1
IASI-NG[11]	MetOp-SG A	2024 (exp)	0.25	0.075
IRS[12]	MTG-S	2024 (exp)	0.625	0.07

bles ($LUTs$[16]) and artificial Neural Networks (NNs[17].).

Both retrieval products are presented as an NH$_3$ profile, from which column amount is derived. However, with each method a fundamental problem persists: the observed radiance spectrum contains insufficient information to accurately determine the vertical distribution of NH$_3$. In the infrared, the absorption spectral signature is non-trivially correlated with the column amount, owing to the contribution of surface-emitted radiation. Hence, methods of this kind are susceptible to inaccuracies of 50%–60% [17], while suffering from a loss of underlying physical information, owing to a focus on statistical correlations. In practice therefore, each retrieval method must assume an approximate, *a priori* profile, in an attempt to constrain this inaccuracy. Again, the high spatio-temporal variability of NH$_3$ distributions compounds this inaccuracy (*cf.* Section 1.2).

Kernel weighting functions (KWFs) are presented as an alternative to both of these methods. The relevant physics and means of construction are presented in Section 2.5 and their performance outlined in Section 3.1. KWFs seek to encompass the additional information required to isolate the effect of the assumed profile on the retrieved column amount. Further, KWFs can be used to interpret the column amount had a different NH$_3$ profile been initially assumed, without repeating the computationally intensive retrieval.

2   Methods

Calculation of an NH$_3$ column amount requires two components: observed radiance line spectra and a subsequent retrieval algorithm. In order to make predictive comparisons between current and future instruments, a radiative transfer model is utilised to produce radiance line


Figure 1: Simplified depiction of atmospheric slab model, illustrating the key principles of a radiative transfer calculation. T_s and T_a denote surface temperature and a general slab temperature, respectively. $B(T)$ denotes the associated spectral density. The red arrows show the path taken by emitted radiation before being detected by the satellite.

spectra. This has the added benefit of allowing for precise control of atmospheric variables (e.g. temperature profile and molecular composition). A basic linear retrieval is chosen for the subsequent calculations.

2.1   Forward Model

The Reference Forward Model[2] (RFM) is a general-use radiative transfer model, originally developed by the AOPP Dept., University of Oxford[18]. The RFM looks to calculate the total radiative transfer to space, summing the surface and atmospheric contributions to give the following expression,

$$R = B(T_s)\mathfrak{T}_s + \int_{\mathfrak{T}_s}^{\mathfrak{T}_\infty} B(T_a)\mathrm{d}\mathfrak{T}_a. \qquad (1)$$

---

[2]I personally operated the RFM program throughout, controlling the desired output by adapting the driver file[18].

A rudimentary atmospheric slab model is presented in order to explain the physical principles underlying this radiative transfer model. Figure 1 demonstrates the basic orientation of Earth's atmosphere. The surface temperature, and a generic slab temperature, are denoted by T_s & T_a, respectively. Both the surface, and each atmospheric slab, emit electromagnetic radiation with a spectral radiance (see section 2.2.1, denoted $B(T_i)$. As the radiation propagates through the atmosphere, its flux is attenuated by a factor known as the transmissibility, $\mathfrak{T} = e^{-\tau}$, phenomenologically described by the Beer–Lambert law. The *optical thickness*, τ, is a dimensionless constant defined by,

$$\tau = \int \kappa \mathrm{d}u, \qquad (2)$$

where u is the partial column amount $(\mathrm{mol/m^2})$ of atmosphere encountered by the radiation, and κ is the absorption cross-section $(\mathrm{m^2/mol})$ and can be taken as constant in our analyses[3]

The total radiative transfer to space, therefore, is the sum of $B(T_s)$ and the integrated of slab contributions, where each slab radiance is attenuated by the optical path length to the atmosphere-space boundary.

Approximating further, the atmosphere is taken to be isothermal – a valid assumption given ammonia predominantly resides in the lower troposphere where the temperature is uniform. Eq (1) simplifies to

$$R = \mathfrak{T}_s B(T_s) + (1 - \mathfrak{T})B(T_a), \qquad (3)$$

where $\mathfrak{T}_\infty = 1$. Taking the atmosphere to be optically thin, we can expand Eq (3) to linear order in u,

$$R = B(T_s) + u\kappa\Delta B, \qquad (4)$$

where we make the approximation: $\mathfrak{T} \approx 1 - u\kappa$, and define $\Delta B = [B(T_a) - B(T_s)]$.

2.2 Radiance Spectra

In order to evaluate the retrieval product for a certain pixel, radiance spectra are needed for atmospheric columns with and without NH_3 present ($\mathbf{y}$ and $\mathbf{y}_0$). However, due to the very low concentration of NH_3 in the atmosphere, great care must be taken to pick a spectral region window in which to observe[19].

---

[3]κ does hold some pressure dependence via *pressure broadening*; however, the maximum Lorentz linewidth ($\sim$0.1 cm^{-1}) associated with this process is unlikely to be resolved by an IASI-like instrument ($\Delta\tilde{\nu}$ =0.5 cm^{-1}).

2.2.1 Planck Function

Planck's law describes the spectral density of electromagnetic radiation emitted by a black body in thermal equilibrium. For our purposes, the inverted form of Planck's law is more useful, with *Brightness Temperature*, T explicitly given as a function of wavenumber, $\tilde{\nu}$, and spectral density B,

$$T(B, \tilde{\nu}) = \frac{hc\tilde{\nu}}{k_B} / \ln[1 + \frac{2h\tilde{\nu}^3}{Bc^2}], \qquad (5)$$

where undefined symbols take their usual meaning.

This formula yields the characteristic black body radiance spectra. By approximating the Earth to be a black body, and with knowledge of atomic absorption spectra for atmospheric species, an appropriate window can be found. Any chosen window must satisfy two properties: (i) a significant and detectable NH_3 absorption feature (ii) spectral signatures of spectator gases must be well modelled, so as to not contribute to the retrieval uncertainty.

The spectral range of 920-940 cm^{-1} was chosen as the focus for project, owing to the dominant absorption feature of NH_3 at 930.75 cm^{-1}, as displayed in figure 2. This absorption feature is molecular in nature, corresponding to the symmetric bending – denoted (ν_2, a_1) – of NH_3, having absorbed radiation satisfying $\tilde{\nu} \approx$ 931 cm^{-1}. The lack of a significant overlapping H_2O feature at this wavenumber minimises the uncertainty associated with the highly variable vertical distribution of H_2O, and thus the second of the aforementioned conditions is also satisfied.

2.3 Linear Retrieval

The premise of a linear retrieval is to map a series of n underlying physical variables to a series of m observables. This relationship can be stated as a matrix equation (generalised dimensions below):

$$\begin{aligned} \mathbf{y} = \mathbf{K}\mathbf{x} \\ (m \times 1) = (m \times n)(n \times 1). \end{aligned} \qquad (6)$$

Following the arguments outlined in Section 1.4, we anticipate a single retrieved quantity, column amount (mol), and so $\mathbf{x}$ is a reduced to a scalar, x, and $\mathbf{K}$ to a column vector of length m, $\mathbf{k}$.

In general, Eq (6) represents a non-linear mapping between our physical parameters and observed radiance. However, in the optically

Figure 2: Individual Brightness Temperatures for four atmospheric gases (H_2O, CO_2, NH_3 and O_3) are plotted across a 20 cm^{-1} wavenumber band. The NH_3 spectrum is re-plotted (purple), but scaled by a factor of 30, to highlight the absorption peak at 931 cm^{-1}. The overall atmospheric spectral pattern is overlaid (dotted, black). All gaseous profiles are taken from the MIPAS dataset[20], and represent a typical, mid-latitude, daytime sounding.

thin regime used to derive equation 4, we assume a linear dependency between column amount and spectral radiance. Hence for small perturbations in x, we can write the respective change in spectral radiance as[4],

$$\mathbf{y} - \mathbf{y}_0 = x\mathbf{k}_0, \quad (7)$$

where we have expanded around $x_0 = 0$, in accordance with the weak absorber limit. The vector, $\mathbf{k}_0$ represents the Jacobian matrix, with elements dictated by the differential response in spectral signature to changes in ammonia column amount:

$$\mathbf{k}_0 = \left.\frac{\partial \mathbf{y}}{\partial x}\right|_{x=x_0}. \quad (8)$$

Finally, the retrieval can be evaluated by inverting Eq (7):

$$x = \mathbf{g}_0(\mathbf{y} - \mathbf{y}_0), \quad (9)$$

where the *Gain Vector*, $\mathbf{g}_0 = (\mathbf{k}_0^T\mathbf{k}_0)^{-1}\mathbf{k}_0^T$.

---

[4]Eq (7) – Eq (9) are adapted from a more generalised derivation found in Rodgers [21].

2.4 Ammonia Profiles

Unlike many atmospheric gases, NH_3 has a very short lifetime[1]. Therefore, any NH_3 produced at the surface will not be convected more than $\sim$15 km into the atmosphere, i.e. just beyond the tropopause. Given the highly variable nature of convection, and the sensitivity of the retrieval to the resulting vertical redistribution of NH_3, different profile shapes must be tested. Hence, I constructed five test profiles, concentrated within the troposphere as depicted in figure 3.
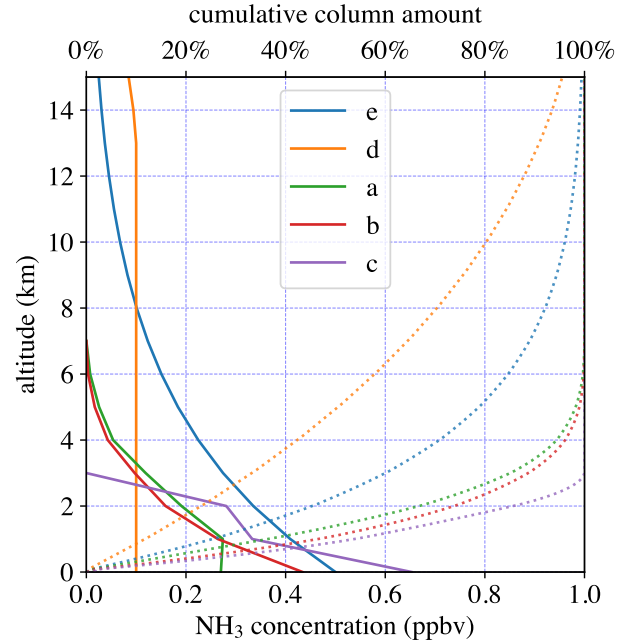

Figure 3: NH_3 test profiles, defined in terms of molar concentration, specifically the volume mixing ratio, v. Profiles are labelled as follows: (e) an exponentially decaying profile; (d) the mid-latitude, daytime profile from the MIPAS project; approximately, it represents a uniform distribution [20] (a) GEOS-CHEM midpolluted site [22]; (b) TES observed, highly polluted site [23] (c) interpolation up to a 4 km mixing level height (MLH)[15]. The cumulative column amount for each profile is overlaid (dotted).

2.4.1 Column Amount Calculation

Given any ammonia profile, the vertical column amount can be calculated, allowing for direct comparison with our spectroscopic retrieval. We first assume hydrostatic balance,

$$\frac{\mathrm{d}p}{\mathrm{d}z} = -g\rho_{\mathrm{m}}M_{\mathrm{air}}, \quad (10)$$

where ρ_{m} is the molar density $(\mathrm{mol/m^3})$ and M_{air} is molar mass of air $(\mathrm{kg/mol})$. Assuming atmospheric layers to be horizontally uniform, we define the total column amount of NH_3 $(\mathrm{mol/m^2})$, U_{NH_3}, in terms of the partial column amounts,

$$U_{NH_3} = \int \mathrm{d}u_{NH_3}. \qquad (11)$$

Introducing the volume mixing ratio, v, via the following relation $\mathrm{d}u_{NH_3} = v\mathrm{d}u_{\mathrm{air}}$, and substituting Eq (10) into Eq (11), we arrive at

$$U_{NH_3} = \frac{1}{gM_{air}} \int v(p)\mathrm{d}p. \qquad (12)$$

A hydrostatic approximation is used to convert our altitude-dependent profiles into pressure-dependent ones and a Simpson's Rule numerical integration scheme is sufficient to compute the column amount to within 1-2% accuracy. It will prove useful to calculate the cumulative column amounts at different altitudes for future analysis. This is depicted in figure 3. The three profiles adapted from from modern literature, {a,b,c}, should be considered more realistic in characterising NH_3 distribution. Note these profiles concentrate as much as 95% of the column amount in the first 6 km. Hence, any further physical analysis can be primarily focused at this altitude range.

2.5 Kernel Weighting Functions

The primary limitations of weak absorber (i.e. NH_3) retrievals is the dimension of the retrieval product, a 1×1 scalar variable, corresponding to total column amount. No information about the vertical distribution of NH_3 is obtained and it will subsequently be shown that the retrieval product itself is highly sensitive to the assumed profile[24]. Kernel weighting functions (KWFs) are a method of dealing with this variability by providing a weighting function, $A(p)$, to be applied to an NH_3 profile of choice. Scaling the NH_3 concentration at each altitude by the associated kernel weight, and subsequently performing a column integration, yields the expected retrieval had the gain matrix, Eq (9), been constructed from the preferred profile, as opposed to an arbitrary, *a priori* profile. The expected column amount for the updated profile is thus,

$$U'_{NH_3} = \frac{1}{gM_{air}} \int A(p)v(p)\mathrm{d}p. \qquad (13)$$

2.5.1 KWF Construction

KWFs represent how the retrieved column amount, U', depends on the assumed vertical profile, $u(p)$ (and hence $v(p)$). Differentiation of Eq (13) yields the following definition,

$$A(p) = \frac{\mathrm{d}U'}{\mathrm{d}u} = \left(\frac{\mathrm{d}U'}{\mathrm{d}R}\right)\left(\frac{\mathrm{d}R}{\mathrm{d}u}\right). \qquad (14)$$

The second equality draws on the mutual relation of U' and u to a single radiance value, R. This definition can be expanded to encompass a spectrum of measured radiance values across our observation window. Defining the observable vector, $\mathbf{y} = \{R(\tilde{\nu}_i)\}^T \equiv \{y_i\}^T$, we can rewrite Eq (14) as

$$A(p) = \sum_{i=0}^{i=m} \left(\frac{\mathrm{d}U'}{\mathrm{d}y_i}\right)\left(\frac{\mathrm{d}y_i}{\mathrm{d}u}\right). \qquad (15)$$

The second bracketed term contains the KWF's pressure-dependence and can be found via differentiation of Eq (4) to give, $\frac{\mathrm{d}y_i}{\mathrm{d}u} = \kappa\Delta B$.

The first term is the underlying Jacobian response, and it proves helpful to restate our linearised retrieval, Eq (7),

$$\mathbf{y} - \mathbf{y}_0 = xU_r\mathbf{k}, \qquad (16)$$

invoking the retrieval quantity, x, where $x := U'/U_r$, and U_r is a reference column amount, specified by the *a priori* profile. Inverting Eq (16), *cf.* Eq (9), we find $\frac{\mathrm{d}U'}{\mathrm{d}y_i} \equiv g_i$, *viz.* components of the *Gain Vector*. Overall, the KWF takes the following form:

$$A(p) = \mathbf{g}\left(\frac{\mathrm{d}\mathbf{y}}{\mathrm{d}u}\right). \qquad (17)$$

In order to evaluate this function, we discretise our space into arbitrary slabs, each with a standard concentration, e.g. $x = 1$ ppbv. The observed spectrum, $\mathbf{y}$, is individually calculated for each slab, until the KWF is smoothly defined across the required altitude range.

3 Results

3.1 Reproduction Matrix

As discussed, the retrieval algorithm is highly sensitive to the *a priori* ammonia profile. The reproduction matrix, displayed in figure 4, demonstrates this with aid of a heatmap to quantify the fractional distance any combination of assumed profile and observed spectra is
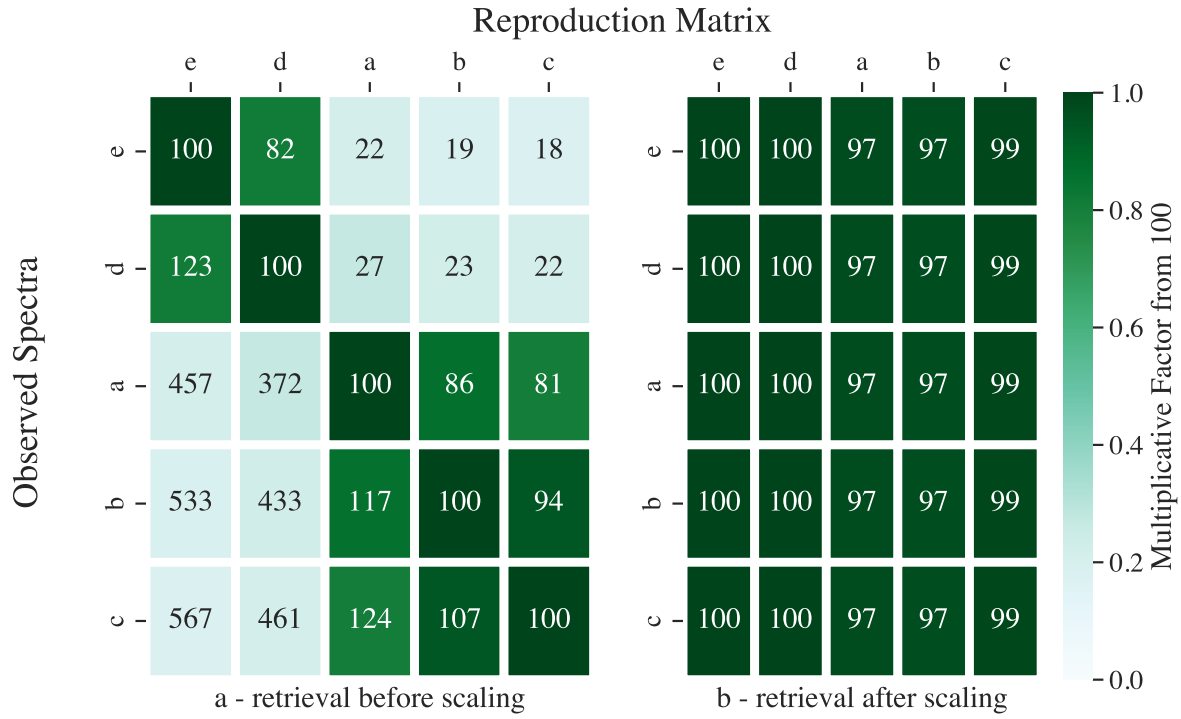
Reproduction Matrix

	e	d	a	b	c
e	100	82	22	19	18
d	123	100	27	23	22
a	457	372	100	86	81
b	533	433	117	100	94
c	567	461	124	107	100

a - retrieval before scaling

	e	d	a	b	c
e	100	100	97	97	99
d	100	100	97	97	99
a	100	100	97	97	99
b	100	100	97	97	99
c	100	100	97	97	99

b - retrieval after scaling

Figure 4: Matrix of retrieval values before (a), and after (b), the kernel weighting functions are applied. The reproduction matrix shows 5×5 combinations of assumed gain matrix, $\mathbf{g}$ (along the top), and observed spectra $\mathbf{y}$ (down the side). The number in each cell is the percentage of the *a priori* column amount that is retrieved for a given observed spectra. The shading depicts the fractional distance a given cell is from retrieving 100%.


Figure 5: The kernel weighting functions for profiles {a,b,c,d,e} are plotted as functions of altitude. A mid-latitude, daytime temperature profile is overlaid(black, dashed) for reference.

from retrieving the exact reference NH_3 column amount. The retrieval is expressed as a percentage in order to normalise for different *a priori* column amounts. Figures 4a & 4b depict the retrieval before and after I applied the relevant kernel weighting functions.

As a check for physical consistency, we test the case where the observed profile matches the internal, *a priori* profile of the retrieval. Trivially, this corresponds to the KWF, $A(p) = 1$, reducing Eq (13) to $U' = U_r$, and thus implying $x = 1$. The diagonal elements of figure 4a are, to within numerical integration error, exactly 100%, thus demonstrating agreement between KWF theory and observed results.

The off-diagonal elements highlight the magnitude of the error for an incorrectly assumed profile. Finally, note the off-diagonal elements are reciprocals of their transpose elements. This is to be expected given the linear nature of the retrieval algorithm: multiplicative scaling factors becoming dividing ones.
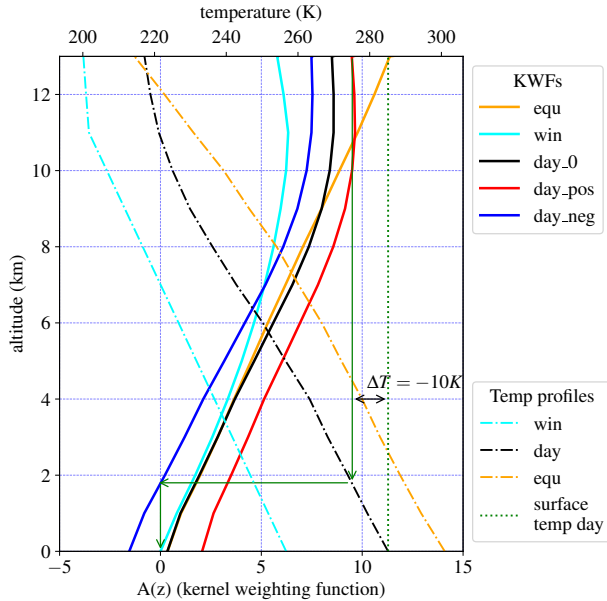
Figure 6: Assuming NH_3 profile c, KWFs are plotted for three temperature contrasts, $TC = \{-10\ K, 0\ K, 10\ K\}$, labelled as 'day-neg', 'day-0' and 'day-pos'. Three further KWFs are shown for $TC = 0$ K but with profiles covering polar, mid-latitude and equatorial regions, labelled as 'win', 'day-0' and 'equ'. The overlaid arrows (green) are used to aid the reader in discerning the KWF value for the 'day-pos' profile, once the air temperature matches the ground temperature.

3.2 KWF Results

I constructed five kernel weighting functions, applying Eq (17) to each of our test profiles. These functions, with a standard temperature profile for reference, are plotted as functions of altitude in figure 5. The temperature response of our weighting curves is immediately obvious. Below ∼10 km, the scale factor linearly increases with height, as the temperature decreases steadily throughout the troposphere. However, each curve displays a clear turning point and subsequent response decay, perfectly in alignment with the temperature inversion at the tropopause. These phenomena are in keeping with Eq (15), where the only altitude variation in A is due to $\frac{dy}{du} \sim \Delta B$, i.e. the difference in surface and atmospheric radiance (and temperature).

So far, KWFs had been constructed assuming both the temperature contrast (TC) and temperature profile (TP) are uniform. To better understand how these two parameters might affect

a KWF, their values were systematically varied for a chosen NH_3 distribution (profile c). Results are displayed in figure 6.

3.2.1 Temperature Contrast Dependency

I chose to construct response curves for three TCs, a range of values encapsulating the majority of routinely observed atmospheres [25]. The effect on KWFs is clear in figure 6: a greater TC corresponds to a greater response at all altitudes. The negative response for $TC = -10$ K, at altitudes < 2 km, is especially noteworthy. This implies that increases in ammonia concentration acts to reduce the atmospheric absorption; alternatively put, *emission-like* features of NH_3 are observed. With the aid of the arrows overlaid in figure 6, it is evident $A(z) = 0$ for z satisfying $T_a = T_s$. Recalling the linearised slab equation, Eq (4), the spectral radiance is equal to a constant, mass-independent value, in the case $T_a = T_s$ and thus $\Delta B = 0$. Hence, our k_i for will be identically zero and all retrievals, and scale factors thereafter, will be zero. The observed result is therefore in keeping with theoretical prediction.

3.2.2 Temperature Profile Dependency

Again, I chose to draw on three temperature profiles, spanning a range of profiles that one might reasonably expect to observe. A strong correlation is observed between scale factor and temperature profile for all three profiles. All three response curves converge at the origin, in keeping with the theoretical expectation for $TC = 0$ K. The three response curves diverge very little for altitudes < 6 km, i.e. the altitude range primarily contributing to the column amount (*cf.* figure 3). Therefore, we deduce the impact of assumed TP is less than that of assumed TC. The KWFs defined by the TC exhibit less variation across the key altitude range, and hence have a lesser effect on the weighted mass path, U'.

4 Retrieval Uncertainty

We evaluate the uncertainty for any given retrieval, following the usual error propagation procedure discussed in Rogers [21]. The uncertainty in column retrieval, σ_x, is related to the error in radiance spectra, $\sigma_{\mathbf{y}}$, via

$$\sigma_x^2 = (\mathbf{k}^T \mathbf{k})^{-1} \sigma_{\mathbf{y}}^2. \tag{18}$$

In arriving at this expression, a lower bound for the uncertainty in each spectral measurement is assumed; it is identically equal to the instrument noise, giving the covariance matrix for $\mathbf{y}$ the following form, $\mathbf{S_y} = \sigma_\mathbf{y}^2 \mathbb{I}_m$.

4.1 Instrument Uncertainty Comparison

Column retrievals were recalculated for the test profiles by independently applying each instrument's spectral resolution, as defined in table 1. Spectral features of a pure NH_3 atmosphere are plotted in figure 7, as viewed by each instrument. Setting $\sigma_\mathbf{y}$ equal to the instrument $NE\Delta T$, retrieval errors are calculated and depicted in figure 8.

It is promising that the anticipated retrieval uncertainty for next generation instruments appears superior to the current tranche of satellites. IASI-NG is set to be particularly impressive, with a factor of 4 improvement in retrieval noise over its predecessor. MTG-IRS will combine low retrieval uncertainty with a greatly improved temporal resolution. The host satellite, MTG-S, will orbit with a 30 minute repeat cycle for full disk coverage (FDC)[12], improving issues of temporal resolution that have previously made short-term tracking of ammonia difficult.
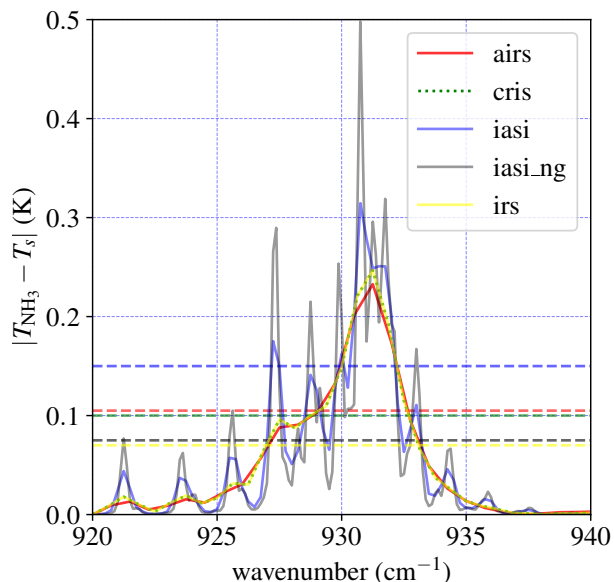

Figure 7: Brightness temperatures for a pure NH_3 atmosphere are calculated for the five specified spectral resolutions, and the absolute difference from the surface temperature is plotted. Instrument $NE\Delta Ts$ are also overlaid (dashed).
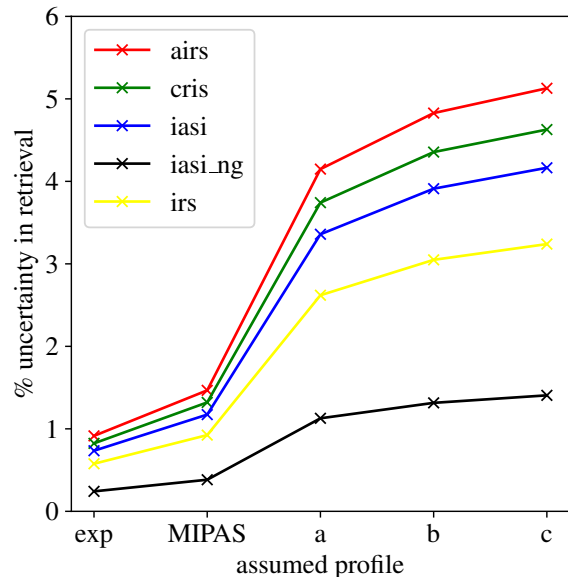

Figure 8: Fractional uncertainty in profile retrieval is expressed as a percentage for each instrument (AIRS - Red, CrIS - Green, IASI - Blue, IASI-NG - Black, IRS - Yellow).

5 Conclusions

5.1 Interpreting Retrieval Reproduction

The primary objective of this project was to demonstrate a) the necessity and b) the efficacy of using KWFs to better interpret the retrieved total column amount. Figure 4a highlights the large variation reference column reproduction, and thus provides a baseline against which we can evaluate the effectiveness of KWFs. The reproduction matrix can be used to quantify the similarity of any two profiles. Simply, the more dissimilar any two profiles, the greater the error in retrieved column amount.

Figure 4b demonstrates the significant improvement in retrieval reproduction when employing KWFs to scale our profiles. Improvements are most substantive with the 'e' and 'd' profiles, while a slight systematic error is introduced for the more physical profiles, 'a', 'b' and 'c'. I believe this discrepancy can be attributed to truncation errors arising from the numerical integration. The latter three profiles are defined in a discrete, piecewise manner which is more susceptive to numerical errors. This is especially true at low altitudes, where the pressure is higher and thus the mass slab contributions are greater, evidenced by figure 3.

5.2 Future Work

I identify two key areas for future work. Firstly, an extensive library of KWFs should be assembled, produced by variation of all relevant atmospheric parameters. Single vector decomposition (SVD) can be applied to this library, and subsequently used to derive a linear mapping between statistical variation in the absorption spectra and atmosphere composition. This would allow the KWFs to be defined by a limited number of coefficients, thus generalising the process of adapting the arbitrarily assumed gain matrix to that defined by an entirely different state vector.

Secondly, KWFs should employed in fast detection methods of NH_3 for real atmospheric data. A comparative analysis in retrieval sensitivity should be conducted to evaluate the efficacy in using KWFs as an alternative to other fast detection techniques being developed, e.g. LUTs and NNs.

References

[1] Van Damme et al. 2021. doi.org/10.1088/1748-9326/abd5e0.

[2] Zhenqi Luo et al. 2022. doi.org/10.5194/acp-22-10375-2022.

[3] Van Damme et al. 2018. doi.org/10.1038/s41586-018-0747-1.

[4] Liu et al. 2010. doi.org/10.1073/pnas.091365810.

[5] Stocker et al. Climate change 2013: The physical science basis. Technical report, 2013.

[6] Tournadre et al. 2020. doi.org/10.5194/amt-13-3923-2020.

[7] Erisman et al. 2007. doi.org/10.1016/j.envpol.2007.06.033.

[8] Dammers et al. 2019. doi.org/10.5194/acp-19-12261-2019.

[9] Bloom. 2001. doi.org/10.1109/IGARSS.2001.976838.

[10] Lambrigtsen et al. 2004. doi.org/10.1109/IGARSS.2004.1370798.

[11] Crevoisier et al. 2014. doi.org/10.5194/amt-7-4367-2014.

[12] Stephen Tjemkes et al. 2015. doi.org/10.1364/HISE.2015.HW1B.2.

[13] Blumstein et al. 2004. doi.org/10.1117/12.560907.

[14] Justice et al. 2002. doi.org/10.1016/S0034-4257(02)00084-6.

[15] Guo et al. 2021. doi.org/10.1029/2020JD033475.

[16] Van Damme et al. 2014. doi.org/10.5194/acp-14-2905-2014.

[17] Whitburn et al. 2016. doi.org/10.1002/2016JD024828.

[18] Dudhia. 2016. doi.org/10.1016/j.jqsrt.2016.06.018.

[19] Dammers et al. 2015. doi.org/10.5194/acp-15-12789-2015.

[20] Endemann. 2005. doi.org/10.1364/FTS.2005.FTuC1.

[21] Rodgers. *INVERSE METHODS FOR ATMOSPHERIC SOUNDING. Theory and Practice.* 2000.

[22] Shephard et al. 2011. doi.org/10.5194/acp-11-10743-2011.

[23] Walker et al. 2012. doi.org/10.5194/acp-12-11213-2012.

[24] Shephard et al. 2011. doi.org/10.5194/acp-11-10743-2011.

[25] Wang et al. 2022. doi.org/10.1016/j.atmosenv.2022.119256.