

# Uncovering Ammonia Hotspots: Resolving Localised Emission Sources from Coarse Satellite Observations

Project Number: AO17

Author: Joel Adams

Supervisor: Dr Anu Dudhia

May 11, 2025

## Abstract

The identification and categorisation of major ammonia emission sources is important for informing environmental policy and mitigating the adverse impacts elevated ammonia concentrations have on both human health and ecosystems. Satellites such as the Infrared Atmospheric Sounding Interferometer (IASI) provide global observations twice daily, but coarse spatial resolution presents difficulties in accurately locating individual emission sources. This study explores the use of linear retrieval techniques in constructing high-resolution maps of ammonia, using IASI data. Two complementary reconstruction methods are developed and evaluated using synthetic data: one based on optimal estimation and another that iteratively updates the distribution. Both methods are shown to successfully locate the synthetic sources at a spatial resolution of approximately 2 km. When applied to real satellite data over the Po Valley for August 2018, the methods identify five significant point sources, highlighting the practical utility of this approach.

## 1 Introduction

Ammonia ( $\text{NH}_3$ ) is primarily released due to agricultural activity, particularly through livestock waste and the use of nitrogen-based fertilisers<sup>[1]</sup>. However, it is a key contributor to the formation of fine particulate matter (PM<sub>2.5</sub>) in the atmosphere through its interaction with other pollutants such as diesel fumes<sup>[1]</sup>. This detrimentally affects air quality, leading to cases of pulmonary and cardiovascular diseases. It is estimated that the inhalation of PM<sub>2.5</sub> as a whole leads to 29,000 – 34,000 additional deaths a year in the UK alone<sup>[2]</sup>. Beyond human health impacts, the deposition of nitrogen as a result of ammonia emissions disrupts the natural chemical balance of ecosystems. Three detrimental processes occur: eutrophication, acidification and leaching<sup>[3]</sup>. Large concentrations of ammonia lead to a build up of nitrogen, a form of eutrophication. Those species better adapted to nitrogen rich environments thrive at the expense of other species, and thus biodiversity is harmed<sup>[4]</sup>. In addition, ammonia undergoes nitrification, reacting with oxygen to release hydrogen ions. This decreases the pH level of the soil, causing acidification, while the negative nitrate ions react with cations essential for plant growth such as calcium ( $\text{Ca}^{2+}$ ), magnesium ( $\text{Mg}^{2+}$ ) and potassium ( $\text{K}^+$ ) and cause soil leaching<sup>[5]</sup>. To mitigate these adverse effects, ammonia distributions must be mapped and major sources found in order to inform policy decisions.

The Infrared Atmospheric Sounding Interferometer (IASI) is a Fourier Transform Spectrometer, operating in the 645 – 2760  $\text{cm}^{-1}$  spectral range, and, since 2007, has been mounted to three MetOp satellites (A, B and C). These satellites operate in polar sun-synchronous orbits, ensuring coverage occurs at the same local time each day, allowing for consistent, twice-daily measurements<sup>[6]</sup>. Prior to the launch of IASI, measurement of atmospheric ammonia was greatly limited to sparse ground-station networks in only certain countries, such as the UK, USA and the Netherlands<sup>[7]</sup>. IASI was the first satellite instrument with sufficient spectral resolution to detect ammonia absorption lines within the infrared region, and this remains the only method of remotely sensing ammonia<sup>[8]</sup>. The

improved spatial and temporal coverage given through remote sensing allows for analysis such as point source identification, the basis of this paper.

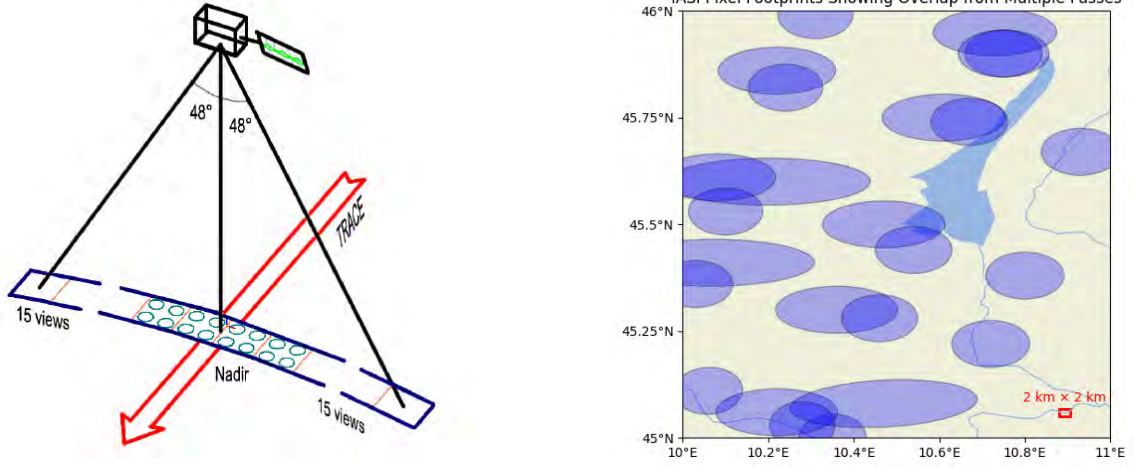


Figure 1: Cross track measurement profile of IASI<sup>[9]</sup> (left) and a small selection of satellite measurements over a  $1^\circ \times 1^\circ$  area, demonstrating the use of overlap to gain sub-pixel information (right). In reality, many more pixels would be needed to construct a high-resolution map. Lake Garda is included for scale reference, with length 52 km, along with a box showing the target high-resolution of  $2 \text{ km} \times 2 \text{ km}$ .

Although each individual pixel is relatively coarse — comprising of ellipses between 12 km and 40 km in length — by using data over an extended time period there is substantial overlap between pixels, as seen in Figure 1b. This spatial oversampling allows information to be derived at a much higher resolution than the raw pixel size.

Previous work completed by Clarisse *et al.*<sup>[10]</sup> have used iterative averaging techniques to construct global high-resolution maps, a technique they label “super-sampling”. They found a large proportion of sources were “animal feeding operations” or industrial emitters. In this paper, the use of linear retrieval methods are used in contrast to this “super-sampling” technique. These models assume the measurement process is a linear operation, and can be specifically tailored to finding point source emitters in a geographically small area, as opposed to the more global outlook of previous studies. In addition, the techniques in this paper are performed with measurements taken over the time scale of only one month, allowing for these small-scale localised maps to be updated on a near-real-time basis. This represents a methodological shift from the global outlook with longer time scales seen in previous studies<sup>[8][10][13]</sup>.

The detection of ammonia in the atmosphere provides a reliable indicator of emission locations due to ammonia’s short atmospheric lifetime, typically on the order of a few hours<sup>[11]</sup>. This means the site of detection is roughly the site of emission. Previous studies, such as the work of Clarisse *et al.*<sup>[10]</sup> incorporate wind transportation models as a necessity to producing the ‘super-sampled’ model. However, in this study wind transportation effects have been neglected as the short atmospheric lifetime, coupled with the assumption that wind direction is sufficiently variable and weak, is such that the effect will be to blur the spatial distribution, as opposed to systematically displacing emission locations. Incorporating wind transportation explicitly in future work could improve the localisation of point sources, particularly in those cases of strong, persistently unidirectional flow.

## 2 Theory

### 2.1 Linear Retrieval

In this paper, a linear retrieval approach is taken to construct a high-resolution ammonia distribution from low-resolution satellite imagery. First, the vector  $\mathbf{y} \in \mathbb{R}^m$  is defined which consists of the

ammonia concentration values taken from satellite measurements, where  $m$  is the number of satellite measurements. The vector  $\mathbf{x} \in \mathbb{R}^n$  is also defined to represent the required high-resolution map of ammonia concentrations. The relationship between the unknown distribution and the known satellite measurements is assumed to be linear—hence linear retrieval—and so is modelled as

$$\mathbf{y} = \mathbf{K}\mathbf{x} \quad (1)$$

where  $\mathbf{K} \in \mathbb{R}^{m \times n}$  is known as the observation operator. Each row of this matrix encodes the spatial extent of the corresponding satellite measurement. The action of  $\mathbf{K}$  on  $\mathbf{x}$  is to average the high-resolution ammonia distribution over the region covered by the respective satellite pixel. Thus, it produces the concentrations as measured by the satellite. The mathematical problem is thus as follows: given a known matrix  $\mathbf{K}$  encoding the spatial distribution of the measurement pixels, find the most probable vector  $\mathbf{x}$  of high-resolution distribution that gives the known vector  $\mathbf{y}$  of ammonia concentrations from satellite measurements.

Due to the number  $m$  of satellite measurements not in general being equal to  $n$ , the observation operator  $\mathbf{K}$  is not square and thus the inverse is not well-defined. Hence, a cost function  $J$  is introduced, to be minimised.

$$J = (\mathbf{y} - \mathbf{K}\mathbf{x})^T \mathbf{S}_y^{-1} (\mathbf{y} - \mathbf{K}\mathbf{x}) \quad (2)$$

Note the inclusion of  $\mathbf{S}_y$ , the covariance matrix of the satellite measurements. This is a diagonal matrix containing the variances of the corresponding satellite measurements. Its inverse weights the measurements inversely according to their variances, ensuring measurements with higher variance contribute less to the cost function. Finally, the derivative of  $J$  is taken with respect to each  $x_i$ .

$$\frac{\partial J}{\partial \mathbf{x}} = -2\mathbf{K}^T \mathbf{S}_y^{-1} (\mathbf{y} - \mathbf{K}\mathbf{x}) = 0 \quad (3)$$

$$\Rightarrow \mathbf{K}^T \mathbf{S}_y^{-1} \mathbf{y} = \mathbf{K}^T \mathbf{S}_y^{-1} \mathbf{K} \mathbf{x} \quad (4)$$

$$\Rightarrow \mathbf{x} = (\mathbf{K}^T \mathbf{S}_y^{-1} \mathbf{K})^{-1} \mathbf{K}^T \mathbf{S}_y^{-1} \mathbf{y} \quad (5)$$

This formulation thus allows for the most probable distribution of  $\mathbf{x}$  to be found, given  $\mathbf{K}$  and  $\mathbf{y}$ <sup>[12]</sup>. Indeed, this approach can be improved immediately with the inclusion of *a priori* knowledge. This can be formalised using Bayes' theorem, and the result is quoted<sup>[12]</sup>:

$$\mathbf{x} = (\mathbf{K}^T \mathbf{S}_y^{-1} \mathbf{K} + \mathbf{S}_x^{-1})^{-1} (\mathbf{K}^T \mathbf{S}_y^{-1} \mathbf{y} + \mathbf{S}_x^{-1} \mathbf{x}_a) \quad (6)$$

Here,  $\mathbf{x}_a$  is defined to hold the *a priori* knowledge of the high-resolution ammonia concentration values, and the matrix  $\mathbf{S}_x$  to be the covariance matrix associated with this knowledge. Importantly, the inclusion of  $\mathbf{S}_x$  within the inverse provides a way to numerically stabilise this inverse. Without retaining  $\mathbf{S}_x$ , were a grid point to be entirely uncovered by satellite measurements, this would result in a column of  $\mathbf{K}$  being entirely zero, rendering the inverse ill-defined. By including a finite variance in the *a priori* estimate, the solution becomes better constrained, and the inverse remains non-singular. It also allows reconstruction to be performed even if the number of satellite measurements ( $m$ ) is less than the number of grid points ( $n$ ), making shorter time periods viable.

Finally, we identify the error covariance matrix as:

$$\mathbf{S} = (\mathbf{K}^T \mathbf{S}_y^{-1} \mathbf{K} + \mathbf{S}_x^{-1})^{-1} \quad (6)$$

This matrix quantifies how both the *a priori* and measurement covariance propagates through the linear retrieval, with the diagonal entries representing the variances for each element of the high-resolution grid.

## 2.2 Spatial Extent of Pixels

The observation operator  $\mathbf{K}$  must now be filled. As has been explained above, each row of the matrix is associated with a single satellite measurement, and performs a mean over the ammonia distribution found within the pixel. Thus, the spatial extent of the satellite pixel must be known. IASI has a field of view of 12km when at nadir (looking directly downwards), and has a swathe of

120 measurements across a maximal angle of  $\pm 48.3^\circ$ <sup>[6]</sup>, as can be seen in Figure 1a. This causes the circular measurement pixel to stretch into an ellipse off nadir, with a maximum major axis of around 40 km. An equation thus needs to be developed to link the angle of measurement from directly upwards (zenith angle  $\mathcal{Z}$ ) to the pixel stretch ( $\ell$ ).




Figure 2: The constructions used to formulate the pixel stretch off nadir. The satellite flies at height  $H = 817$  km above the Earth surface. It is assumed the Earth is perfectly spherical with a radius  $R = 6378$  km and that on the scale of one pixel the Earth can be treated as flat. The equations are solved for  $\ell$  and thus the pixel size is found, given the zenith angle  $\mathcal{Z}$ .

Note  $\varphi$  is the angular pixel size as taken from the satellite and is therefore constant. At nadir, it is easy to identify  $\varphi = \frac{12}{H}$ . Secondly, the sine rule is used on both diagrams as shown in Figure 2:

$$\frac{\ell}{\sin(\varphi)} = \frac{z}{\sin(90 + \mathcal{Z} - \varphi)} \quad \frac{\sin(\mathcal{Z})}{R + H} = \frac{\sin(\Theta)}{R} \quad (7)$$

and finally, the cosine rule is used on Figure 2 (left) to find

$$z^2 = R^2 + (R + H)^2 - 2R(R + H) \cos(\mathcal{Z} - \Theta) \quad (8)$$

Therefore, 3 equations have been defined for the 3 unknowns  $\Theta$ ,  $z$  and  $\ell$  and so the pixel size  $\ell$  can be found, in km. By including the latitude of measurement, along with the angle of orbit, a major axis  $u$  and minor axis  $v$  can be determined, expressed in degrees latitude and longitude. This spatial extent then allows for the determination of which high-resolution grid-points in  $\mathbf{x}$  lie within the pixel. The associated entries in the observation operator  $\mathbf{K}$  are set to  $1/N$ , with  $N$  being the number of high-resolution points covered by that pixel. This ensures a mean of all concentration values within the pixel is taken, analogous to the satellite measurement being taken.

## 3 Method

### 3.1 Simple Averaged Method (SAM)

In this paper, three models are used of increasing sophistication to enable both comparison and validation. The first model, serving as a baseline, is the simplest approach possible. It assigns to each high-resolution grid point an ammonia concentration determined by taking the simple average over all satellite measurements covering that location. While this method provides a coarse representation of the underlying distribution, it suffers significantly from blurring, as would be expected given its crude nature.

### 3.2 Optimal Estimate Solution (OES)

A more sophisticated approach is to use the theory developed in Section 2. Here, there is assumed to be no prior knowledge of the underlying distribution, and thus  $\mathbf{x}_a = 0$ . The covariance matrix

$\mathbf{S}_x$  is retained within the inverse:

$$(\mathbf{K}^T \mathbf{S}_y^{-1} \mathbf{K} + \mathbf{S}_x^{-1})^{-1} \quad (9)$$

This provides numerical stabilisation, as discussed at the end of Section 2. Optimal performance was found when the diagonals of  $\mathbf{S}_x$  were set to  $100\times$  the mean of the diagonals of  $\mathbf{S}_y$ , corresponding to an *a priori* standard deviation of  $10\times$  the mean satellite measurement standard deviation.

Furthermore, the off diagonal entries of  $\mathbf{S}_x$  encode *a priori* knowledge of correlations between points within the high-resolution grid. More concretely,  $[\mathbf{S}_x]_{ij}$  is defined to be equal to the *a priori* covariance between grid-point  $i$  and grid-point  $j$ . In this formulation, an anti-correlation is introduced between adjacent points, which encourages the constructed distribution to localise within individual grid points, and contrast more clearly from the surrounding noise. This anti-correlation is implemented by assigning negative values to the off-diagonal entries corresponding to adjacent points, set here at 10% of the diagonal entries of  $\mathbf{S}_x$ . Consequently, when a particular grid point indicates a high ammonia concentration, adjacent points are constrained to lower values. This makes point sources more distinguishable against surrounding noise, and reduces the blurring of point sources during the construction process.

### 3.3 Iterative *a priori* Refinement (IAPR)

A final method is used to further enhance the detection of point sources, and proceeds as follows. First, each grid point is retrieved individually under the assumption ammonia is only found at that point, using Equation 6 with  $\mathbf{K}$  reduced to the column vector corresponding to the grid point being tested, and  $\mathbf{x}$  reduced to a scalar. For each grid point, the cost function  $J$  is then found to test how well the actual distribution is simulated by the single point. By selecting the point associated with the lowest cost function, the best estimate of the distribution is found, were the distribution to be localised at a single point.

This value then becomes the *a priori* knowledge for the next iteration. Each unassigned grid point is then tested again in a similar way, but with  $\mathbf{x}_a$  now holding the previously assigned point, with all other points kept with  $\mathbf{x}_a = 0$ . This allows the second point to be fitted given the knowledge of the other, retrieved grid point. The third iteration is then completed with knowledge of the previous two points, and thus the distribution is progressively constructed, one grid point at a time. Throughout the process, both  $\mathbf{S}_y$  and  $\mathbf{S}_x$  remain constant, with  $\mathbf{S}_x$  once again utilised to ensure anti-correlation between adjacent grid-points.

Due to the ‘greedy’ nature of this approach, the algorithm is expected to prioritise grid points most likely to contain high ammonia concentrations first, as these grid points individually provide the closest approximation to the complete measured distribution when considered in isolation. Therefore, this method excels at rapidly identifying point sources. Indeed, the iterative process can be terminated before all grid points are retrieved, allowing for early stoppage once a sufficient representation of the distribution is found. The cost function reduces most rapidly as any point sources are identified, and slows significantly as the noisy background distribution is retrieved.

## 4 Validation and Testing

With the three methods to be used established, the models were tested for their ability to reconstruct point sources effectively and reliably. For this, a synthetic high-resolution distribution vector  $\mathbf{x}_0$  was created, which consisted of two point sources of size 2km across, as can be seen in Figure 3 (top left). Though the distribution was synthetic, both the satellite measurement’s location and footprint were taken from a real IASI dataset, and used to create the observation operator  $\mathbf{K}$ . The dataset corresponded to a month’s worth of measurements over August 2018 in a  $1^\circ \times 1^\circ$  area located at  $45^\circ\text{N}, 10^\circ\text{E}$ , and consisted of 429 individual measurements. This location includes Lake Garda (length 52km) and thus provides a length scale for comparison (as seen in Figure 3). Thus, synthetic satellite measurements were created by applying  $\mathbf{K}$  to  $\mathbf{x}_0$ . Finally, Gaussian noise was added to the synthetic satellite measurements, with a standard deviation of 10% of the maximum signal amplitude. This ensures all measurements receive a non-zero level of synthesised noise, consistent with instrument uncertainty. The methods of reconstruction as described in Section 3



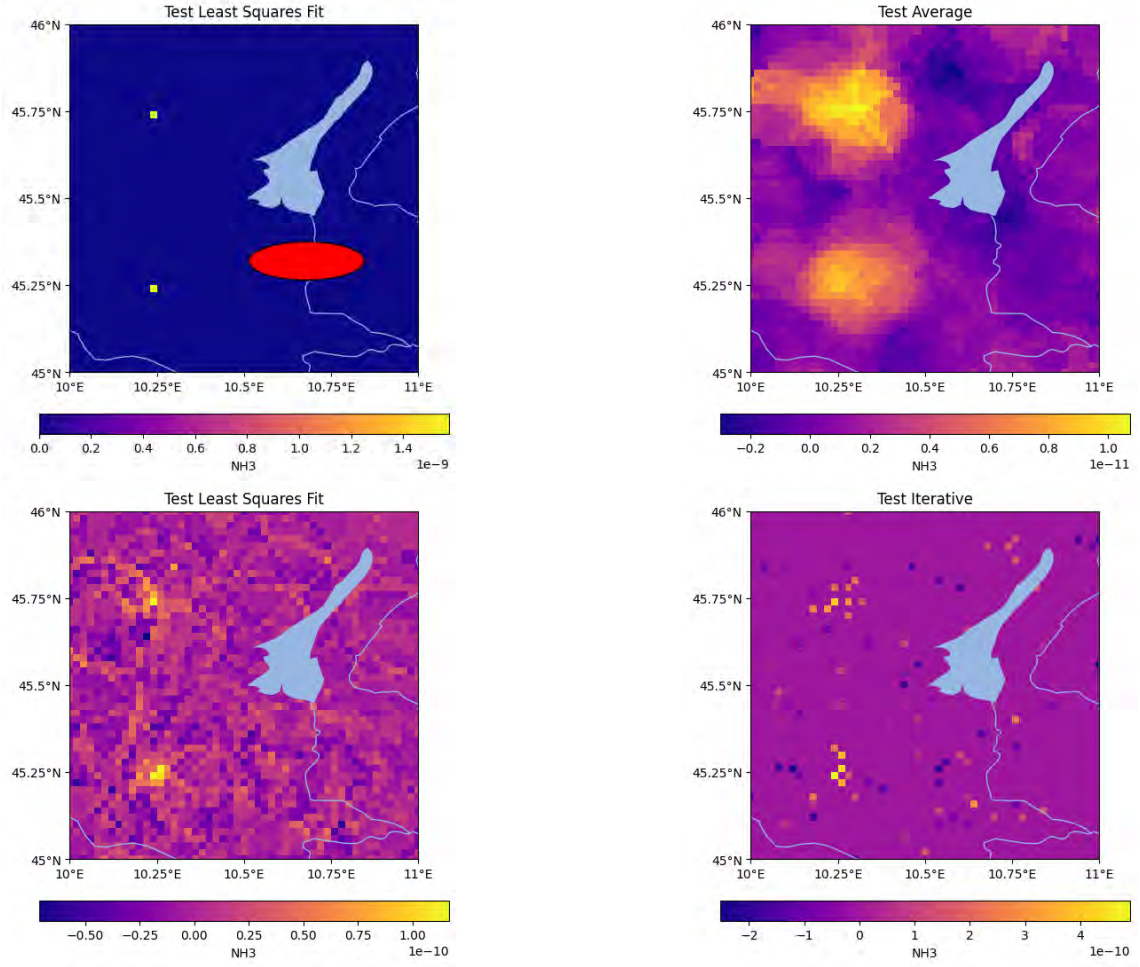


Figure 3: The original synthetic distribution in yellow with a representative satellite measurement footprint in red for comparison (top left) and the three methods used to reconstruct the distribution, SAM (top right), OES (bottom left) and finally IAPR (bottom right).

were then applied to reconstruct the ammonia distribution as best as possible, and the results are plotted in Figure 3.

It is clear that the SAM offers little insight into the exact location of the point sources; however, it provides a coarse representation of the overall distribution at low computational cost. As a result, it is well suited for surveying large areas of land to identify regions best suited for the more refined methods to be employed. Both the OES and IAPR methods are able to accurately and precisely locate point sources, enabling their exact locations to be determined. When applied to real data, the use of both methods in parallel will serve as a means of cross validation, enhancing confidence in the constructions. With the recollection that the footprint of IASI has a minimum diameter of 12 km<sup>[6]</sup>, and the size comparison of source to satellite footprint displayed in Figure 3, the significant improvement to locating a point within 2 km is evident.

## 5 Results

With the models developed having been tested on synthetic data, a real distribution over the Po Valley was used to attempt the identification of point sources. This is the larger dataset from which the measurements used for validation and testing were pulled from, and was derived using the Oxford MORSE retrieval algorithm applied to data from IASI-A and IASI-B, August 2018 (A Dudhia, pers. comm., generated for this project). The Po Valley is the area of Europe

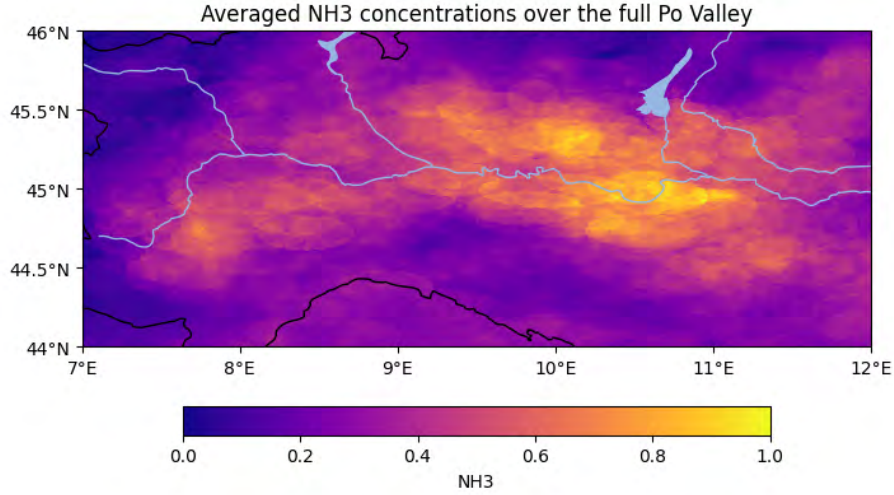


Figure 4: A large scale reconstruction using the simple average method. It is clear the ammonia is concentrated around the route of the Po river, following the majority of agriculture and industry as expected. Note the values have been rescaled as this method does not accurately reproduce absolute concentrations.

which consistently displays the highest concentration of atmospheric ammonia<sup>[14]</sup>, and thus it was deemed most likely to display significant, identifiable sources. The month of August was similarly chosen as seasonal cycles are shown to yield maximum ammonia concentrations through spring and summer<sup>[14]</sup>. The dataset was initially filtered to remove data points with too high cloud cover, as this obscures measurements and would detract from the results<sup>[7]</sup>.

The covariance matrix of the satellite measurements  $\mathbf{S}_y$  was populated with the associated standard deviations as derived within the MORSE retrieval algorithm. These represent the propagated measurement noise, and reflect the noise-limited nature of the IASI retrieval. The uncertainties primarily depend upon the strength of emitted radiance, and are thus driven by surface and atmospheric temperature<sup>[15]</sup>. An analysis of the signal-to-noise ratio (SNR), defined as the ratio of absolute ammonia concentration to its associated standard deviation for each measurement, revealed the most reliable measurements were concentrated in the central region of highest ammonia emission — roughly 9.5° to 11.5°E and 44.5° to 45.5°N — as seen in Figure 4. Within this region, a mean SNR of 3.14 was found over 1143 individual measurements — indicating a moderate but usable level of confidence in the retrieval values.

Initially, the SAM was used across the entire Po Valley to gain an understanding of the rough distribution, as seen in Figure 5. The initial distribution looks as expected, with the majority of ammonia lying along the route of the river, as this is where the majority of agriculture and industry is located. The North and West have significantly lower levels, as the Alps precludes much farming or factories.

With the rough distribution thus mapped, a more rigorous identification of point sources was attempted using the methodologies outlined above. An analysis of the propagated covariance matrix  $\mathbf{S}$  over the domain 9.5° to 11.5°E and 44.5° to 45.5°N at a resolution of 0.02°, revealed a mean SNR after construction of just 0.37, and a maximum of 2.03. This low value indicates that, on average, the retrieved distribution is dominated by noise. However, this is to be expected for a noise-limited measurement process such as that of IASI, and likely reflects the genuinely low ammonia concentrations across much of the domain. Taking only those points with an SNR greater than 1.4 yields 48 individual points, and comparing these with the 48 points of highest retrieved ammonia concentrations reveals high similarities, see Figure 5. This overlap suggests the strongest ammonia signals are also associated with the most confidently retrieved values, increasing the likelihood that these regions reflect genuine emission sources rather than noise artefacts.

With areas of high-concentration and relatively low noise identified, further investigation into each of these was conducted. This involved taking smaller domains and performing the IAPR

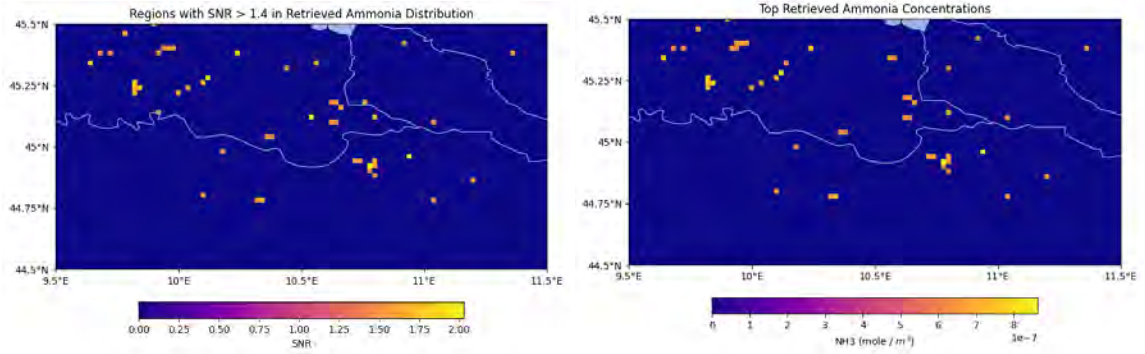


Figure 5: A comparison of the retrieved points with signal to noise ratio of greater than 1.4 (left) to the corresponding number of highest ammonia concentration (right). The strong spatial similarity provides an indication that areas of high-concentration are more likely to correspond to genuine emission sources.

method to ensure the high-concentration points remain identifiable and spatially localised. Secondly, domain sensitivity was tested by shifting the retrieval domain by small amounts and ensuring sources persist. Taken together, these tests indicated five points of interest, all of which correspond to identified areas in Figure 5. The process undertaken to find these points is illustrated in Figure 6, and a large scale map identifying the sources is to be found in Figure 7.

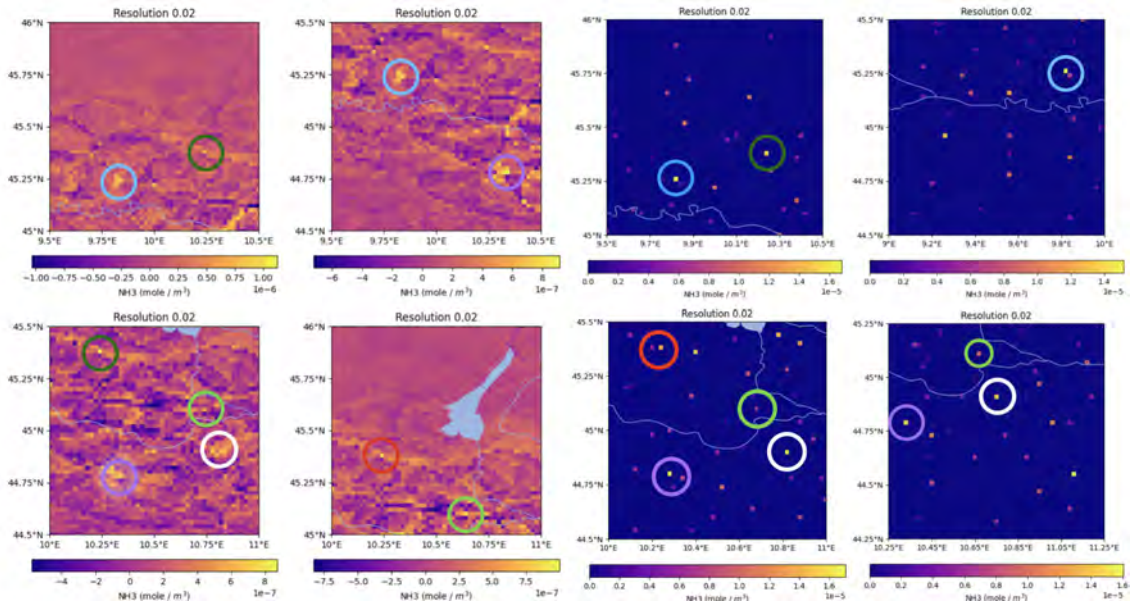


Figure 6: An illustration of the process of identifying the point sources. Five point sources are identified to be consistent through a range of plots. Colours used correspond to these in Figure 7.

An inspection of satellite imagery over the indicated areas reveals points 1 through 4 lie in intensively farmed agricultural regions and there are several candidate farm complexes within the predicted areas which could plausibly be responsible. To establish a more definitive causal link, further investigations are necessary. According to existing literature<sup>[13]</sup>, agricultural hotspots were consistently found to be “intensive animal farming” operations. Thus, a further investigation into the specific nature of these farms, and the existence of such intensive animal farming in the specified areas, would provide a definitive link.

Point 5, located in a predominately urban area in the south suburbs of Parma, presents a more difficult case. It could be speculated that this is due to industrial activity, but without further evidence, a definitive conclusion is impossible to draw.



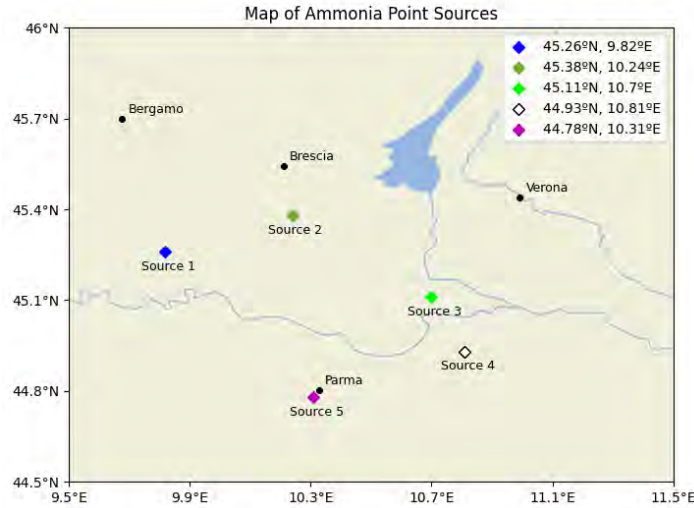


Figure 7: A map of ammonia point sources detected over the Po Valley. Sources 1 through 3 appear to be agricultural sources, with source 4 potentially industrial. However, without further analysis, definitive conclusions are hard to draw.

An attempt was made to separate the data from MetOp-A with that from MetOp-B to enable an independent comparison between the datasets as a form of validation. However, the process of halving the dataset led to a substantial decrease in retrieval quality: the maximum SNR was just 1.50 for MetOp-B. This led to difficulties in identifying the emission points from surrounding noise, and thus the results were inconclusive. Nonetheless, Source 4 remained prominent in both datasets, providing some additional support for its identification as a true emission point. To utilise this validation technique, longer observation periods would be required, a point future studies may wish to consider.

## 6 Conclusion and Further Work

Two linear retrieval methods, and a simple averaging method, have been evaluated for their suitability for the identification of point sources of ammonia, using data from IASI. Synthetic distributions were created and used to test these retrieval methods and it was demonstrated that both linear retrieval techniques were capable of producing reconstructions of the distribution to an accuracy of around 2km, provided the point sources were of sufficient amplitude when compared to the background distribution. One month's worth of data was chosen, as this allows the map of point sources to be updated on a frequent time frame.

The true measurements of ammonia concentrations were then analysed across the Po Valley, a known area of high ammonia concentration in Europe. It was found that the retrieval methods suffered with a low signal to noise ratio across much of the domain. However, there were areas of higher concentrations of ammonia, and these had comparatively better signal to noise ratios. Of these areas, five presented themselves as sufficiently identifiable over a range of more empirical tests, such as domain shifting. These were identified as potential ammonia emission sources.

Further investigation into these five sources through satellite imagery did not yield clear evidence of specific activities directly responsible for the elevated ammonia levels. While no definitive explanation was immediately obvious it is hypothesised further research into specific land use and more comprehensive local knowledge could provide further evidence. Although the comparison of MetOp-A measurements to those from MetOp-B proved challenging due to the low dataset size, it was found at least one of these sources was identifiable in both independent datasets. Further work with longer observation windows would allow this technique to be expanded upon, leading to increased validation. Moreover, incorporating more subtle effects into the models such as horizontal transport due to winds could enhance the accuracy and may lead to better identification of clear emitters.

## References

- [1] Wyer, K.E. *et al.* (2022) ‘Ammonia emissions from agriculture and their contribution to fine particulate matter: A review of implications for human health’, *Journal of Environmental Management*, 323, p. 116285. Available at: <https://doi.org/10.1016/j.jenvman.2022.116285>.
- [2] Marais, E.A. *et al.* (2023) ‘The health burden of air pollution in the UK: a modelling study using updated exposure-risk associations’, *Lancet* (London, England), 402 Suppl 1, p. S66. Available at: [https://doi.org/10.1016/S0140-6736\(23\)02099-8](https://doi.org/10.1016/S0140-6736(23)02099-8).
- [3] Guthrie, S. *et al.* (2018) ‘The impact of ammonia emissions from agriculture on biodiversity’, The Royal Society. Available at: [https://www.rand.org/pubs/research\\_reports/RR2695.html](https://www.rand.org/pubs/research_reports/RR2695.html)
- [4] Emmett, B.A. *et al.* (2011) ‘Interpretation of evidence of nitrogen impacts on vegetation in relation to UK biodiversity objectives’, JNCC Report, 449. Available at: <https://data.jncc.gov.uk/data/2bef3139-5e46-44dc-b6e4-d0881bfea900/JNCC-Report-449-FINAL-WEB.pdf>
- [5] Cui, L. *et al.* (2021) ‘Effects of Nitrification Inhibitors on Soil Nitrification and Ammonia Volatilization in Three Soils with Different pH’, *Agronomy*, 11(8), p. 1674. Available at: <https://doi.org/10.3390/agronomy11081674>.
- [6] EUMETSAT and European Organisation For The Exploitation Of Meteorological Satellites (2022) ‘IASPCS1C0100IASI Principal Components Scores Fundamental Data Record Release 1 - Metop-A and -B’. European Organisation for the Exploitation of Meteorological Satellites. Available at: [https://doi.org/10.15770/EUM\\_SEC\\_CLM\\_0084](https://doi.org/10.15770/EUM_SEC_CLM_0084).
- [7] Van Damme, M. *et al.* (2015) ‘Towards validation of ammonia (NH<sub>3</sub>) measurements from the IASI satellite’, *Atmospheric Measurement Techniques*, 8, pp. 1575–1591. Available at: <https://doi.org/10.5194/amt-8-1575-2015>.
- [8] Van Damme, M. *et al.* (2014) ‘Global distributions, time series and error characterization of atmospheric ammonia (NH<sub>3</sub>) from IASI satellite observations’, *Atmospheric Chemistry and Physics*, 14. Available at: <https://doi.org/10.5194/acp-14-2905-2014>.
- [9] Clerbaux, C. *et al.*: Monitoring of atmospheric composition using the thermal infrared IASI/MetOp sounder, *Atmos. Chem. Phys.*, 9, 6041–6054, <https://doi.org/10.5194/acp-9-6041-2009>, 2009.
- [10] Clarisse *et al.* ‘Tracking down Global NH<sub>3</sub> Point Sources with Wind-Adjusted Superresolution’. *Atmospheric Measurement Techniques* 12, no. 10 (17 October 2019): 5457–73. Available at: <https://doi.org/10.5194/amt-12-5457-2019>.
- [11] Xie, Y. *et al.* (2024) ‘NH<sub>3</sub> Emissions and Lifetime Estimated by Satellite Observations with Differential Evolution Algorithm’, *Atmosphere*, 15(3), p. 251. Available at: <https://doi.org/10.3390/atmos15030251>.
- [12] Rodgers, C. D. (Clive D.) 2000, *Inverse methods for atmospheric sounding : theory and practice*, World Scientific, [River Edge, N.J.], Singapore
- [13] Van Damme, M. *et al.* (2018) ‘Industrial and agricultural ammonia point sources exposed’, *Nature*, 564(7734), pp. 99–103. Available at: <https://doi.org/10.1038/s41586-018-0747-1>.
- [14] Van Damme, M. *et al.* (2022) ‘On the weekly cycle of atmospheric ammonia over European agricultural hotspots’, *Scientific Reports*, 12(1), p. 12327. Available at: <https://doi.org/10.1038/s41598-022-15836-w>.
- [15] Piera Raspollini *et al.* (2014) ‘Comparison of the MIPAS products obtained by four different level 2 processors’, *Annals of Geophysics*, 56, p. 9. Available at: <https://doi.org/10.4401/ag-6338>.