

Retrieving tropospheric refractivity structure using interferometry of aircraft radio transmissions

Ollie Lewis^{1,a}, Chris Brunt¹, Malcolm Kitchen¹, Neill E. Bowler², and Edmund K. Stone²

¹Department of Physics and Astronomy, University of Exeter, Exeter, EX4 4QL, United Kingdom

^aNow at: National Center for Atmospheric Research, 3450 Mitchell Lane, Boulder, Colorado, United States of America

²Met Office, FitzRoy Road, Exeter, EX1 3PB, United Kingdom

Correspondence: Ollie Lewis (olewis@ucar.edu)

Abstract. Detailed measurements of atmospheric humidity in the lower atmosphere are currently difficult and expensive to obtain. For this reason, there is interest in the development of low-cost, high-volume opportunistic technologies to acquire measurements of tropospheric humidity. We demonstrate the use of interferometry to measure the atmospheric refraction of the Automatic Dependent Surveillance-Broadcast (ADS-B) radio transmission routinely broadcast by commercial aircraft.

5 Atmospheric refraction is strongly influenced by changes in humidity, and refractivity observations have proved to be an effective source of humidity information for numerical weather prediction models. A prototype ADS-B interferometer has been developed that can simultaneously perform angle-of-arrival (AoA) interferometry and decode ADS-B signals. Combining the measured AoA of the ADS-B signal with the known position of the aircraft (information contained within the ADS-B signal) allows the bending of the signal due to refraction to be determined. Combining the measured bending of numerous ADS-B
10 signals allows for information concerning the refractivity structure to be extracted. An adjoint model was derived and used to retrieve synthetic one-dimensional refractivity profiles in a variety of atmospheric conditions. The results from an experiment using a prototype ADS-B interferometer are shown and initial refractivity profiles are retrieved. Sources of uncertainty in the observations and the retrieved refractivity profiles were explored and future work was suggested.

Copyright statement. The works published in this journal are distributed under the Creative Commons Attribution 4.0 License. This li-
15 cence does not affect the Crown copyright work, which is re-usable under the Open Government Licence (OGL). The Creative Commons Attribution 4.0 Licence and the OGL are interoperable and do not conflict with, reduce or limit each other.

© Crown copyright 2025

1 Introduction

The global distribution of water vapour partially governs the transport of heat and the radiation budget of the atmosphere
20 (Trenberth et al., 2009; Held and Soden, 2000). In the lower atmosphere, the spatial and temporal variability in water vapour strongly influences the formation and intensity of extreme weather events through condensation and subsequent release of latent heat (Trenberth et al., 2005). The high sensitivity of atmospheric processes to water vapour presents the requirement

for accurate observations for both climatological and meteorological purposes. Despite its importance in climate and driving extreme weather events, the spatial and temporal variability of tropospheric humidity is difficult to capture adequately using existing observing systems (Rieckh et al., 2018). Radiosondes provide the principal source of in situ tropospheric water vapour observations, capable of retrieving high resolution vertical profiles of humidity. However, the limited spatial distribution of launch sites and low frequency of radiosonde ascents (typically once daily) substantially restrict the ability to retrieve the full four-dimensional structure of water vapour in the lower atmosphere. Another source of in situ humidity observations is provided by the Aircraft Meteorological Data Relay (AMDAR). The onboard water vapour sensing system (WVSS) instrument can provide detailed profiles of humidity, particularly as the aircraft traverses the moist lower atmosphere (Petersen et al., 2016). However, only a fraction of AMDAR reports currently contain usable humidity data, with the majority restricted to North America (Petersen et al., 2016; Ingleby et al., 2019; Pauley and Ingleby, 2022).

Developments in remote sensing of atmospheric water vapour have substantially improved the ability of observing systems to capture the variability and distribution of tropospheric humidity. Concerning remote sensing of humidity, surface-based retrievals using lidars can provide excellent temporal coverage of the variability in humidity throughout the extent of the lower atmosphere directly above the sensor (Guo et al., 2023). However, individual units are expensive and limited in number, restricting the distribution of retrieved vertical profiles of tropospheric humidity for use in numerical weather prediction (NWP) (Gaffard et al., 2021). The maturation of satellite-based remote sensing of humidity has dramatically increased the global coverage of humidity observations. Satellite observations, such as those from passive microwave and infrared-sounding satellites, have proved to be a crucial source of indirect humidity observations. In particular, microwave humidity-sounding has had a substantial positive impact on short-range forecasting skill (Geer et al., 2017). Microwave humidity-sounding satellite observations are an especially crucial source of humidity information over the data-sparse open oceans, where the positive impact of radiance-derived humidity observations in these regions has been demonstrated to propagate downstream with the atmospheric flow over much of Europe and landmasses in the southern hemisphere. However, the large weighting functions (2-3 km) used to quantify the vertical distribution of humidity severely limit the resolution of retrieved profiles (Rieckh et al., 2018; Wang et al., 2024) relative to the vertical resolution of NWP models such as the Met Office United Kingdom (UK) Variable-resolution (UKV) model (average vertical spacing in the troposphere is on the order ~ 100 m) (Tang et al., 2013).

High resolution direct/indirect measurements of humidity are challenging to obtain. However, the high spatial and temporal variability of water vapour in the lower atmosphere results in significant changes in the optical properties of the atmosphere. A valuable source of indirect humidity information for use in NWP is provided through observations of refractivity. An example would be the Global Navigation Satellite System (GNSS) radio occultation technique (Kursinski et al., 1997; Kuo et al., 2000; Healy and Thépaut, 2006), which exploits the refractive bending of radio signals between a low Earth orbit (LEO) satellite and a GNSS satellite to retrieve the refractive index (refractivity) structure of the atmosphere. Although GNSS radio occultation can provide accurate high resolution vertical profiles of refractivity, the retrievals can be significantly impacted by strong gradients in refractivity (Xie et al., 2006). Refractivity observations have also been obtained using radar phase shifts (Fabry et al., 1997; Weckwerth et al., 2005) and zenith tropospheric delay (ZTD) measurements of GNSS satellite radio signals (Rohm, 2012; Zhao

et al., 2019). However, the enormous spatial and temporal variability in refractivity make accurate, high resolution retrievals challenging.

There is an increasing interest to acquire atmospheric observations using commercial aircraft, which are mandated to broadcast digitally encoded three-dimensional position via the Automatic Dependent Surveillance-Broadcast (ADS-B) system for air traffic control and safety purposes (ICAO, 2012). Meteorological parameters derived using ADS-B and Mode Selective Enhanced Surveillance (Mode-S EHS) data include temperature (de Haan, 2011; Stone and Kitchen, 2015; Mirza et al., 2019) and wind (de Haan et al., 2013; Mirza et al., 2016) data. Aircraft-derived observations have been exploited by the UK Met Office (Stone and Pearce, 2016; Stone, 2018; Li, 2021) and Royal Netherlands Meteorological Institute (KNMI) (de Haan and Stoffelen, 2012; de Haan et al., 2013) to retrieve high density wind data. The potential to retrieve atmospheric refractivity data using aircraft-derived observations has recently been explored (Lewis et al., 2023a). A two-element interferometer was used to measure the bending of refracted ADS-B signals and analogies with GNSS radio occultation were investigated (Lewis et al., 2023b). In this study we demonstrate the potential to retrieve one-dimensional refractivity profiles using synthetic and real observations obtained using a prototype ADS-B interferometer.

2 Methods

2.1 Interferometry of radio transmissions

Atmospheric refraction of ADS-B radio signals transmitted by aircraft can be detected using vertically-orientated, two-element interferometers. The total bending of the radio signal can be determined by combining the interferometrically observed angle-of-arrival (AoA) with the known point of origin of the signals (information encoded within the ADS-B transmission). The observed AoA measurement technique is described in detail in Lewis et al. (2023a). Measurable bending only occurs at grazing incidence ($\text{AoA} \lesssim 2^\circ$), therefore a clear view of the horizon is required to acquire meteorologically useful observations. A measurable phase difference between the signal received at the upper antenna and the signal received at the lower antenna is induced by the path length excess (geometry depicted in Fig. 1). The phase difference, ϕ , can be written in terms of the path excess as

$$\phi = 2\pi \frac{\mathbf{B} \cdot \hat{\mathbf{z}}}{\lambda} = 2\pi \frac{B \sin(\beta)}{\lambda}, \quad (1)$$

where \mathbf{B} is the vertical baseline vector (displacement between the upper and lower antennas), $\hat{\mathbf{z}}$ is the unit vector in the direction of the received ADS-B signal, β is the observed AoA (elevation angle of the received signal) and λ is the signal wavelength (~ 0.275 m for a 1090 MHz transmission). An Ettus Universal Software Radio Peripherals (USRP) X300 software-defined radio (SDR) was used to determine the phase difference by computing the conjugate product of the signals from each interferometer element.




Figure 1. Geometry of the ADS-B interferometry observing system.

2.2 Atmospheric refraction

The refractive properties of the atmosphere can be described using refractivity, which is the parts-per-million (ppm) excess of the refractive index above the unity vacuum value. The total refractivity at a given point in the atmosphere can be decomposed into the dry refractivity (contribution of the dry atmospheric gases to the total refractivity, N_{dry}) and the wet refractivity
90 (contribution of water vapour to the total refractivity, N_{wet}) as

$$N_{\text{dry}} = q_1 \frac{P}{T}, \quad (2)$$

$$N_{\text{wet}} = q_2 \frac{e}{T^2}, \quad (3)$$

where q_1 and q_2 are empirically-derived constants given by $q_1 = 77.6 \text{ K hPa}^{-1}$ and $q_2 = 3.73 \times 10^5 \text{ K}^2 \text{ hPa}^{-1}$ respectively (Smith and Weintraub, 1953). The air pressure, temperature and partial pressure of water vapour are given by P [hPa], T
95 [K] and e [hPa] respectively. Other more accurate formulae exist for the calculation of refractivity (e.g. Aparicio and Laroche (2011)) but the simple form is still commonly used. The contribution of water vapour to the total refractivity is generally around 20% near the surface (Möller, 2017) but can fluctuate rapidly over short spatial and temporal scales. Refractivity data was obtained using the 12:00 Coordinated Universal Time (UTC) launch radiosonde data from Watnall, Nottinghamshire (UK) (World Meteorological Organisation (WMO) station number 03354) on the 22 September 2023. The radiosonde data used
100 in this study are available from the University of Wyoming, United States of America. Figure 2 shows the tropospheric dry refractivity profile, as well as the modulation due to the wet refractivity contribution. Above an altitude of around 7 km, the contribution of the wet refractivity to the total refractivity is negligible.




Figure 2. The total refractivity (solid black line) and dry refractivity (solid green line) versus altitude (in kilometres) retrieved using the 12:00 UTC launch radiosonde data from Watnall on the 22 September 2023.

To model the propagation of a radio signal through an atmosphere with a radially-directed refractive index gradient, we apply ray-tracing techniques used in geometric optics. The ray path in a spherically symmetric atmosphere can be defined using Snell's law

$$n \cdot (a + h) \cdot \cos(\epsilon) = \text{constant}, \quad (4)$$

where n is the refractive index, h is the height of the ray above the surface, ϵ is the local elevation angle of the ray and a is the radius of the Earth. The ray propagation can be formulated as an initial value problem, dependent on the initial position and direction of the ray. Here we use the second-order ordinary differential equation (SODE) derived by Zeng et al. (2014) to simulate the propagation. The SODE model is derived using the initial assumption that the height, h , is dependent on the along-beam range, r . Therefore

$$n[h(r)] \cdot (a + h(r)) \cdot \cos(\epsilon) = \text{constant}. \quad (5)$$

Differentiating (5) with respect to r allows the second-order differential equation to be derived (the full derivation is described in detail in Zeng et al. (2014))

$$\frac{d^2 h}{dr^2} = \left(1 - \left(\frac{dh}{dr} \right)^2 \right) \left(\frac{1}{n} \frac{dn}{dh} + \frac{1}{a + h} \right). \quad (6)$$

The introduction of the variable u allows the SODE to be separated into a system of first-order differential equations

$$\frac{dh}{dr} = u, \quad (7)$$

$$\frac{du}{dr} = (1 - u^2) \left(\frac{1}{n} \frac{dn}{dh} + \frac{1}{a + h} \right), \quad (8)$$

which can be integrated over the domain $[0, r_f]$ (where r_f defines the endpoint of the ray) to recover the ray path. We choose to
120 integrate equations (7) and (8) using a third-order Runge-Kutta scheme. The distance travelled by the ray at each step, i , along
the surface of the Earth, s , is given by

$$s_{i+1} = s_i + \sin^{-1} \left(\frac{\cos(\epsilon_i) \Delta r}{a + h_{i+1}} \right) a, \quad (9)$$

where

$$\epsilon_i = \sin^{-1}(u_i), \quad (10)$$

125 is the local elevation angle of the ray at each step. The initial height and direction of the ray are required to model the refraction
of a radio signal through the atmosphere. The altitude of the aircraft can be determined, however, the initial emission angle of
the radio signal is unknown. The refraction is therefore modelled in a time-reversed frame from the position of the receiver to
the reporting aircraft. A prior (first-guess) refractivity profile is used. The squared difference between the altitude of the model
ray after termination and the reported height of the aircraft is used to compute a penalty function. The best estimate of the true
130 refractivity profile is determined by minimising the penalty function. We explore this optimisation problem in more detail in
Section 2.3.

2.3 Optimisation of the refractive index profile

The optimisation problem can be formulated as

$$\min_n \quad J = \sum_{i=0}^{N_B} j_i = \sum_{i=0}^{N_B} ((h_f)_i - (h_t)_i)^2, \quad (11a)$$

$$135 \quad \text{subject to} \quad \dot{h}(r; n, h_0, u_0) - u = 0, \quad \forall r \in [0, r_f], \quad (11b)$$

$$\dot{u}(r; n, h_0, u_0) - (1 - u^2) \left(\frac{1}{n} \frac{dn}{dh} + \frac{1}{a + h} \right) = 0, \quad \forall r \in [0, r_f], \quad (11c)$$

$$h(0; n, h_0, u_0) = h_0, \quad (11d)$$

$$u(0; n, h_0, u_0) = u_0, \quad (11e)$$




Figure 3. Illustration of the single-ray optimisation problem. A ray with initial position, h_0 , and elevation (exaggerated in the diagram for illustrative purposes), β , is traced through an initial-guess refractive index profile until it reaches the aircraft distance, s_t . The penalty function for a single ray is calculated as the squared difference between the ray endpoint, h_f , and the target (aircraft) position, h_t . The intermediate height, distance and local elevation angle of the ray along the path are given by h_i , s_i and ϵ_i respectively.

140 where h and u describe the height (position) and (sine) direction of the ray respectively and \dot{h} and \dot{u} are their respective derivatives with respect to r . The initial position and direction of the ray are given by h_0 and β respectively. J is the total penalty function describing the sum of the individual penalty functions, j_i , corresponding to the squared difference between the end height of the i^{th} ray, $(h_f)_i$, and the i^{th} target height, $(h_t)_i$ ($i = 0, 1, 2, \dots, N_B$, where N_B is the number of broadcasts used in the optimisation problem). In the case of ADS-B interferometry, the target position will be the reported position of the

145 aircraft.

Figure 3 shows the geometry of the system to be optimised for a single ray. The refractive index profile, \mathbf{n} , was optimised such that a ray traced through the profile with an initial position and direction given by h_0 and β (where $u_0 = \sin(\beta)$) respectively will terminate at the target (aircraft) position, h_t . The gradient of the penalty function, J , with respect to the refractive index profile, \mathbf{n} , was required to minimise J via gradient descent. The gradient of the individual penalty function for a single

150 ray can be approximated by computing finite differences over the elements of the refractive index profile $\mathbf{n} \in \mathbb{R}^M$ (where M is the number of refractive index values along the ray trajectory) as

$$\frac{dj[h_f(\mathbf{n})]}{dn_k} \approx \frac{j[h_f(\mathbf{n} + \delta n \mathbf{e}_k)] - j[h_f(\mathbf{n} - \delta n \mathbf{e}_k)]}{2\delta n_k}, \quad (12)$$

where n_k is the k^{th} element of the refractive index profile. The k^{th} element of the vector \mathbf{e}_k ($\mathbf{e}_k \in \mathbb{R}^M$) is unity and zero otherwise and δn is a small increment to the k^{th} element of \mathbf{n} . Each gradient evaluation with respect to a single element of

155 the refractive index profile requires integrating equations (11b) and (11c) over the entire ray trajectory to determine the change in the end position of the ray, h_f . The penalty gradient calculation using finite differences for a single ray is expensive and quickly becomes impractical when evaluating the total penalty function gradient for N_B rays when N_B is large. In a previous investigation by Lewis et al. (2023a), the total penalty function given by (11a) was minimised using simulated annealing, a stochastic global search optimisation algorithm (Kirkpatrick et al., 1983). Although simulated annealing is an effective method
160 for optimising high-dimensional problems with potentially many local minima, the algorithm is very computationally expensive and requires thousands of iterations to converge on a good solution. For this reason, the optimisation procedure using the adjoint state method (Zou et al., 1997; Gill et al., 1981) was used to reduce the computational cost of retrieving the refractivity profile. Applying the adjoint state method to our original optimisation problem, the constrained optimisation problem described by (11) can be converted into an unconstrained optimisation problem using Lagrange multipliers

$$165 \quad \mathcal{L} = j - \int_0^{r_f} \psi (\dot{h} - u) dr - \int_0^{r_f} \mu \left(\dot{u} - (1 - u^2) \left(\frac{1}{n} \frac{dn}{dh} + \frac{1}{a+h} \right) \right) dr, \quad (13)$$

where ψ and μ are the adjoint variables associated with the state variables h and u respectively. The integration of equations (11b) and (11c) over the domain $[0, r_f]$ (where r_f parameterises the endpoint of the ray) defines the trajectory of the ray. The minimisation of (13) corresponds to finding the stationary points of the Lagrangian with respect to the trajectory defined by h and u . Therefore any variations in the Lagrangian with respect to the trajectory are equal to zero (that is, $\delta \mathcal{L}|_h = \delta \mathcal{L}|_u = 0$).
170 The adjoint variables enforce this constraint. Solving for the stationary points of the Lagrangian with respect to h and u we find

$$\psi|_{r=r_f} = \frac{\partial j}{\partial h}, \quad (14)$$

$$\mu|_{r=r_f} = \frac{\partial j}{\partial u}, \quad (15)$$

$$\frac{d\psi}{dr} = \mu(1 - u^2) \left(\left(\frac{1}{n} \frac{dn}{dh} \right)^2 - \frac{1}{n} \frac{d^2n}{dh^2} + \frac{1}{(a+h)^2} \right), \quad (16)$$

$$175 \quad \frac{d\mu}{dr} = 2\mu u \left(\frac{1}{n} \frac{dn}{dh} + \frac{1}{a+h} \right) - \psi. \quad (17)$$

The full derivation is described in Appendix A. The initial conditions of the adjoint variables ψ and μ are defined at the endpoint of the ray trajectory, $h|_{r=r_f}$. The adjoint equations are therefore evolved in reverse from the endpoint to the initial position of the ray. In the simulated retrievals, an atmosphere with spherical symmetry was assumed and values were computed using a

linear interpolation scheme, hence $d^2n/dh^2 = 0$. Since the refractive index generally decreases exponentially with height, the
180 interpolation was performed using the natural logarithm of the refractive index values. Expressing the height-dependent natural
logarithm of the refractive index as $\tilde{n}(h) = \ln(n(h))$, equations (16) and (17) can then be written as

$$\frac{d\psi}{dr} = \mu(1 - u^2) \left(\left(\frac{d\tilde{n}}{dh} \right)^2 + \frac{1}{(a + h)^2} \right), \quad (18)$$

$$\frac{d\mu}{dr} = 2\mu u \left(\frac{d\tilde{n}}{dh} + \frac{1}{a + h} \right) - \psi. \quad (19)$$

The integration of the adjoint variables requires knowing the spatial gradient of the refractive index at the query points h
185 along the ray trajectory. The refractive index profile was optimised using the two-stage optimisation procedure described by
Teh et al. (2022): forward tracing and backward tracing. The forward tracing stage involves integrating (11b) and (11c) using
initial conditions h_0 and u_0 and evaluating the endpoint of the ray trajectory $h|_{r=r_f}$. The backward tracing stage then involves
initialising the adjoint variables (14) and (15) at the termination position and integrating (18) and (19) in reverse along the
same ray trajectory as in the forward tracing stage. The variation of the Lagrangian with respect to the refractive index is then

$$190 \quad \delta\mathcal{L}|_{\tilde{n}} = \int_0^{r_f} \mu(1 - u^2) \left(\frac{d(\delta\tilde{n})}{dh} \right) dr, \quad (20)$$

which does not require an expensive Jacobian matrix solve to evaluate. The quantity $\delta\tilde{n}$ is defined in terms of a linear interpo-
lation scheme as

$$\tilde{n}(h, \mathbf{n}) = \mathbf{w}(h) \cdot \ln(\mathbf{n}) = \begin{bmatrix} w_0(h) & w_1(h) \end{bmatrix} \cdot \begin{bmatrix} \ln(n_0) \\ \ln(n_1) \end{bmatrix}, \quad (21)$$

where n_0 and n_1 are the refractive index values on the underlying grid directly below and above the sample point h respec-
195 tively. The weights $w_0(h)$ and $w_1(h)$ are inversely related to the vertical distance from the refractive index values n_0 and n_1
respectively. Therefore at each step in the integration of the adjoint variables the gradient of the Lagrangian with respect to the
refractive index values on the underlying grid are calculated as

$$\frac{d\mathcal{L}}{d(\ln(n_0))} = \int_0^{r_f} \mu(1 - u^2) \frac{dw_0}{dh} dr, \quad (22a)$$

$$\frac{d\mathcal{L}}{d(\ln(n_1))} = \int_0^{r_f} \mu(1 - u^2) \frac{dw_1}{dh} dr, \quad (22b)$$




Figure 4. The geometric ray path intersecting the refractive index values $\ln(n_0)$ and $\ln(n_1)$ on the underlying grid.

where the relation $\delta \tilde{n} = \mathbf{w} \cdot d(\ln(\mathbf{n}))$ is employed. The geometry of the optimisation scheme at each step in the integration is shown in Fig. 4.

The refractive index field was adjusted iteratively using a gradient descent scheme using a variation of the Adam optimiser (Kingma and Ba, 2014). The minimisation algorithm was terminated after a specified number of iteration steps. The gradient of the penalty function calculated using the adjoint state method was verified using finite differences over the elements of the refractive index profile, as shown in Fig. 5. The penalty function is highly sensitive to the surface refractive index value (as shown in the left panels of Fig. 5), highlighting the importance of accurately knowing the surface conditions in order to constrain the optimised solution. The right panels of Fig. 5 show the magnitude of the relative difference between the gradients calculated using finite differences and the adjoint state method for 1 m and 100 m ray integration step sizes respectively. The value of the gradients computed using finite differences and the adjoint state method converge as the ray step size decreases, however, the computational cost of integrating rays with a very small step size increases significantly. For this study the 100 m integration step size was used as the gradient calculation was sufficiently accurate without being prohibitively expensive to compute for thousands of rays.

In summary, there are two principal observations required to retrieve atmospheric refractivity structure: the observed AoA of the radio broadcast (required to calculate the initial ray direction, u_0) and the position of the aircraft (required to determine s_t and h_t). In the examples above we initialise the adjoint variable μ using the partial derivative of the penalty function with respect to the ray direction u . However, in practice, we have no knowledge of the emission direction of the ADS-B signal. Therefore we define the penalty function only in terms of the respective positions of the ray endpoint and the aircraft position. The adjoint term μ is therefore initialised at the termination position as $\mu|_{r=r_t} = 0$.

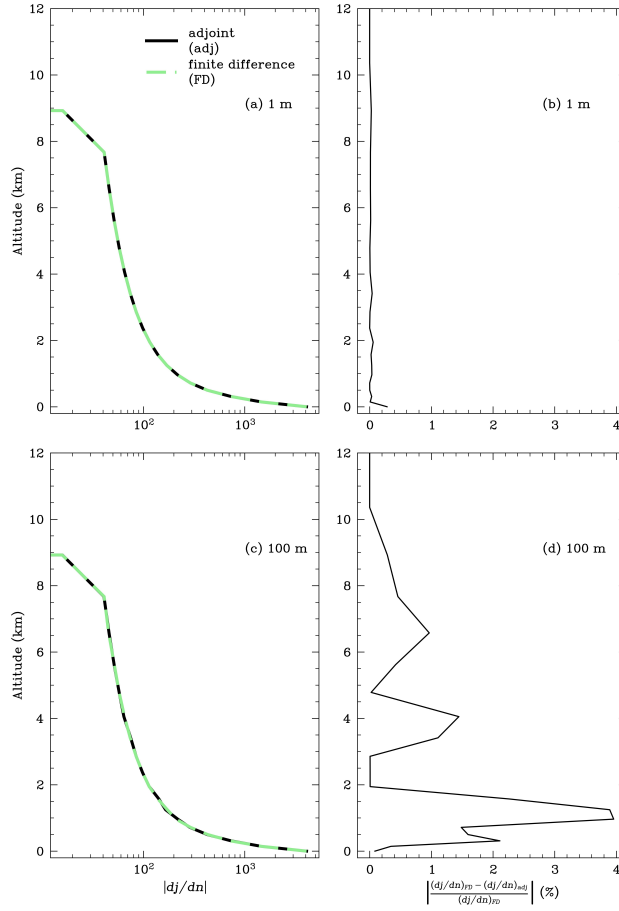


Figure 5. Left panels: The gradient of the penalty function with respect to the refractive index at each altitude point calculated using finite differences (green) and the adjoint state method (black) for **(a)** 1 m and **(c)** 100 m ray integration step size. The penalty function is highly sensitive to the refractive index values close to the surface. **Right panels:** The magnitude of the relative difference between the gradients calculated using finite differences (FD) and the adjoint state method (adj) for **(b)** 1 m and **(d)** 100 m ray integration step size. The ray had an observed AoA of 0.2° . The altitude and distance of the source aircraft were 8.8 km and 350 km respectively.

2.4 Experimental arrangement

220 The observations were obtained using a prototype ADS-B interferometer installed on the Cleve Hill weather radar tower in Shropshire, UK. The experimental setup is shown in Fig. 6. The vertical baseline of the interferometer was 13.86 m, and each of the two elements consisted of a stack of four Yagi-Uda antennas separated by 0.66 m. The stack reduced the beamwidth of each element, helping to mitigate the effects of signals reflected from nearby topography (multipath). The antenna stacks were connected to an Ettus USRP X300 SDR.




Figure 6. The interferometer mounted to the side of the Cleve Hill weather radar tower.

225 Approximately 263,500 refracted ADS-B transmissions were received between an AoA of 0.0° to 2.0° over an ~ 80 -minute period between 08:36 and 09:58 UTC on the 22 September 2023. The synthetic and real retrievals were obtained using the central 10° azimuthal sector of received transmissions. The distribution of all the received transmissions and those used for retrievals are shown in Fig. 7a and 7b respectively. The distribution in reported aircraft altitude within the chosen sector is shown in Fig. 7c. The north-east sector was chosen due to a suitable distribution in reported altitudes for the lower atmosphere
230 to be adequately sampled. The nearby Watnall radiosonde station also allowed for a reference refractivity ground truth to be determined to compare retrieved refractivity profiles to.

The adjoint inversion algorithm was tested on synthetic data generated with ADS-B data recorded using the prototype interferometer at Cleve Hill, Shropshire. The observations used in the analysis were restricted to the central 10° azimuthal sector of the entire observation dataset. A subset of 5000 ADS-B transmissions were chosen randomly from within the central
235 azimuthal sector to use in the synthetic refractivity retrieval experiments. Radiosonde data from the nearby Watnall station in Nottinghamshire (position shown in Fig. 7a) was used to generate a refractivity profile to simulate the refraction of the subset of ADS-B transmissions. Three radiosonde soundings were used to generate the synthetic observations - 12:00 UTC 22 September 2023 (see Fig. 2), 12:00 UTC 18 July 2022 (during an intense heatwave) and 12:00 UTC 15 December 2022

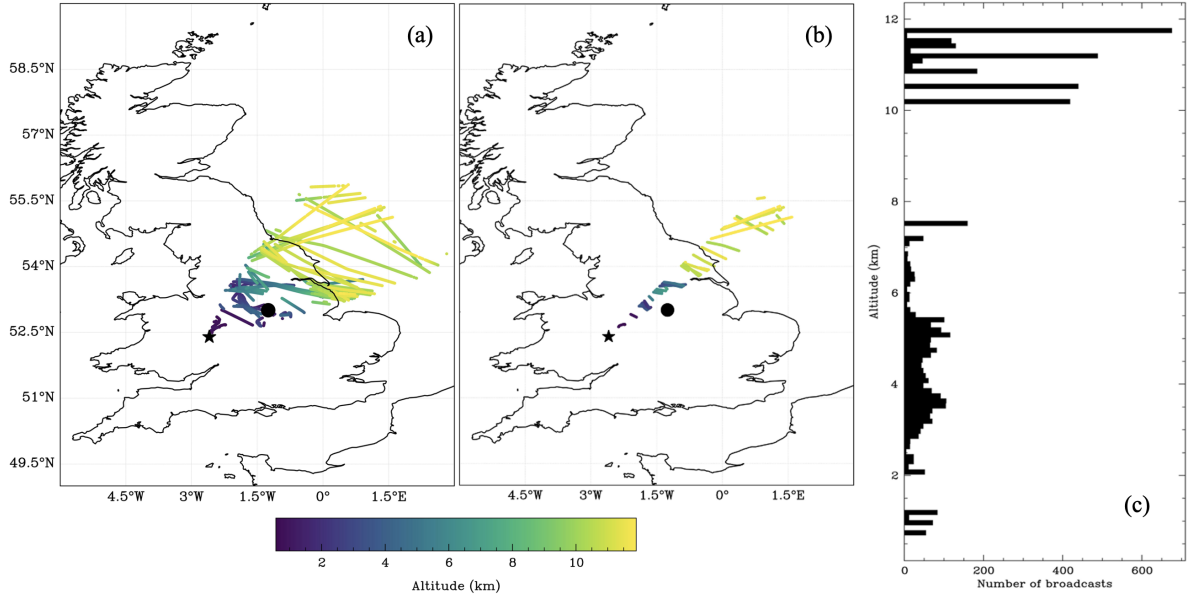


Figure 7. (a) The spatial distribution of all the received aircraft transmissions colour-coded by reported altitude (km). (b) The spatial distribution of the analysed aircraft transmissions colour-coded by reported altitude (km). The positions of the interferometer and the Watnall radiosonde station are indicated by the black star and black circle respectively. (c) Histogram of the distribution of the analysed broadcasts by reported altitude (km). The analysed aircraft transmissions were sampled from the central 10° azimuthal sector of received transmissions on the 22 September 2023.

(during a cold spell). The synthetic observations (heights of synthetic aircraft) were generated by tracing the 5000 rays out from the receiver to the reported distances, s_t , (see Fig. 3) of each of the ADS-B transmissions using the refractivity profiles determined using the radiosonde data. The interferometrically-measured observed AoAs were used as the initial directions of the rays and the height of the ray endpoint was used as the "true" height of the synthetic ray origin (i.e. equivalent to the height of a synthetic aircraft). The method is equivalent to that shown in Fig. 3, except the heights of the synthetic aircraft are given by h_t . The heights of the synthetic aircraft are treated as the target heights h_t in the synthetic retrievals.

The observed AoAs used in the synthetic refractivity retrievals were injected with varying degrees of noise to model the impact of experimental uncertainties on the retrieved refractivity profiles. A prior (first-guess) refractivity profile, given by

$$N(h) = N_0 \exp\left(-\frac{h}{H}\right), \quad (23)$$

was used to initialise the penalty function, where N_0 is the refractivity at the position of the receiver (this was fixed to the radiosonde value during the optimisation procedure) and H is the scale height in kilometres. Rays were traced backwards from the receiver position to the reported distance of the aircraft using the known observed AoAs, as shown in Fig. 3. The total penalty function was then evaluated by summing together the penalty function for each ray (i.e. the squared difference

between the traced ray end height and the height of the synthetic aircraft). Since multipath contamination of the observed AoA measurement is thought to constitute the principal source of uncertainty in the retrieved refractivity profile, three observed AoA measurement noise test cases were investigated. The retrieval algorithm was run with a zero measurement noise case and
 255 two normally distributed noise test cases (with standard deviations of 0.01° and 0.05° respectively).

To account for the ellipsoid shape of the Earth in the one-dimensional ray tracing model, the mean local radius of the ellipsoid at the location of the receiver is used as the effective radius of the Earth. The effective radius is approximately the radius of curvature of the ellipsoid at the location of the receiver, which ensures that contours of constant refractivity are approximately parallel to the curved surface of the ellipsoid. The height above the ellipsoid reported in the ADS-B transmission was left
 260 unchanged. The coordinate transform is described in Appendix B.

3 Results and analysis

3.1 Synthetic refractivity retrievals

The retrieved refractivity profile consisted of 30 levels on a logarithmically-spaced grid between an altitude of 0.575 km (height of receiver above the ellipsoid Earth model, estimated using a GNSS receiver on site) to an altitude of ~ 13 km. A ray step
 265 size of $\Delta r = 0.1$ km was used during the integration. A learning rate of 10^{-7} was used during the optimisation procedure. In order to better constrain the retrieved solution (particularly in the presence of noise) the minimum refractivity after every iteration during the optimisation procedure was fixed to be greater than or equal to the dry refractivity (i.e. assuming pressure and temperature are known and the relative humidity can never drop below 0%). The root mean squared error (RMSE) in the synthetic retrieved refractivity profiles after convergence for the test cases using measurement noises of 0.00° , 0.01° and 0.05°
 270 are shown in Table 1. The RMSE was computed as

$$\text{RMSE} = \sqrt{\frac{\sum_{i=1}^{30} (N_{\text{retrieve},i} - N_{\text{radiosonde},i})^2}{30}}, \quad (24)$$

where N_{retrieve} is the retrieved refractivity and $N_{\text{radiosonde}}$ is the refractivity determined from radiosonde data.

The RMSE in the retrieved profile increases as the standard deviation in the observed AoA noise increases. However, all retrieved profiles show improvement relative to the prior refractivity profile. A synthetic retrieval with no observed AoA
 275 measurement noise was also performed assuming the height of the receiver above the ellipsoid had an error of ± 5 m and ± 10 m. The synthetic retrieval assuming a receiver height error still showed significant improvement in the lower relative to the initial refractivity profile. However, there was little improvement above an altitude of ~ 10 km. The weakly constrained refractivity retrievals at high altitudes are explored in Section 4.

Figure 8a shows the synthetic retrieved vertical refractivity profile for the different observed AoA measurement noise cases
 280 using the 22 September radiosonde data. The 0.00° RMSE profile was obtained from the 5000 interferometrically-measured, unperturbed observed AoAs. The 0.01° and the 0.05° RMSE profiles were obtained by perturbing the 5000 interferometrically-measured observed AoAs according to the corresponding normal distributions. The initial refractivity profile was defined using

Table 1. The RMSE values for the initial refractivity profile and the synthetic retrieved refractivity profile (relative to the reference radiosonde profile). The three angular values refer to the standard deviation of the added noise to the observed AoA.

| Radiosonde data (12:00 UTC) | Initial RMSE (ppm) | Retrieved RMSE (ppm) | | |
|--------------------------------|--------------------|----------------------|-------|-------|
| | | 0.00° | 0.01° | 0.05° |
| 22 September 2023 | 8.21 | 0.76 | 1.42 | 3.11 |
| 18 July 2022 | 9.63 | 1.89 | 3.42 | 7.09 |
| 15 December 2022 | 4.84 | 0.70 | 0.88 | 1.18 |

(23) with a scale height $H = 8$ km. The percentage difference between the retrieved and the radiosonde-derived refractivity profile is shown in Fig. 8b. The 0.05° RMSE profile shows a lower refractivity error standard deviation at higher altitudes than the 0.01° RMSE profile. The increased standard deviation in the AoA measurement error results in a greater number of AoAs being perturbed below 0° for the 0.05° RMSE case. Since the retrievals can currently only be performed using AoAs $\geq 0^\circ$, these are then rejected, resulting in a bias towards perturbing the AoAs towards higher elevation angles. This results in the gradient in the retrieved refractivity profile becoming steeper which, as a consequence of the clamping of minimum refractivity, results in the 0.05° RMSE profiles clustering together with a lower variance. However, the larger RMSE in the observed AoA results in a larger bias in the retrieved refractivity profile.

The synthetic refractivity retrievals were then repeated for two other radiosonde profiles - the 12:00 UTC Watnall soundings for the 18 July 2022 and 15 December 2022 respectively. The synthetic retrievals were performed using synthetic observations of aircraft heights generated from the two respective refractivity profiles in a process identical to that using the 22 September data. The synthetic retrievals using injected measurement noise with a standard deviation of 0.01° and 0.05° both showed improvements relative to the initial guess refractivity profile in the lower atmosphere for the 18 July and 15 December profiles, as shown in Table 1. The synthetic refractivity retrievals demonstrate the ability to obtain useful atmospheric data using measurements of refracted ADS-B transmissions in the presence of random noise.

3.2 Initial retrievals using observations

The retrieval algorithm was then tested on the observational data collected using the prototype ADS-B interferometer. The retrieved refractivity profiles were constrained by fixing the minimum refractivity to the dry refractivity. It is hoped that the technique will be able to resolve variations in humidity structure occurring on timescales of order an hour. Accordingly, the observed AoA measurements from one hour of the experiment were divided into four 15-minute periods between 08:36 and 09:36 UTC and used to retrieve refractivity profiles. The observed AoA ranged between 0.0° and 2.0° and the azimuthal angle range was chosen as the central 10° of the sector shown in Fig. 7. The four observational periods contained 9700 (08:36-08:51 UTC), 8315 (08:51-09:06 UTC), 8185 (09:06-09:21 UTC) and 7589 (09:21-09:36 UTC) received ADS-B transmissions respectively within the sector. Figure 9 shows the:




Figure 8. (a) The synthetic retrieved refractivity profile and the reference 22 September 2023 Watnall radiosonde profile for different synthetic error distributions injected into the observed AoA measurement. (b) The percentage difference between the synthetic retrieved refractivity profile and the reference radiosonde profile for different synthetic error distributions injected into the observed AoA measurement. The shaded regions indicate the limits in uncertainty in the synthetic retrieved refractivity profile assuming no AoA measurement noise and a ± 5 m (red) and ± 10 m (blue) receiver altitude error respectively. The fractional error standard deviation in the retrieved refractivity profiles using twenty different observed AoA noise realisations are given by the blue and red dashed lines for the noise distributions with standard deviations of 0.01° and 0.05° respectively.

- AoA_{LoS} : The line-of-sight (LoS) AoA determined from the reported position of the aircraft;
- $\text{AoA}_{\text{init_model}}$: The synthetic LoS AoA determined from the end position of the ray traced through the initial refractivity profile;
- $\text{AoA}_{\text{ret_model}}$: The synthetic LoS AoA determined from the end position of the ray traced through the retrieved refractivity profile;

as a function of azimuthal angle for the four 15-minute periods of observational data. The LoS AoA refers to the angle above the local horizon of the straight line joining the receiver and the reported aircraft position/end position of a traced ray (the LoS AoA computation is explained in more detail in Lewis et al. (2023a)).

315 The penalty function is minimised when the difference between the AoA_{LoS} and the AoA_{ret_model} is minimised. The AoA_{init_model} and AoA_{ret_model} simulated using the model refractivity profiles (as shown by the blue and green points in Fig. 9 respectively) provides a direct comparison with observable quantities (the AoA_{LoS} determined using the received ADS-B transmission). The AoA_{init_model} and AoA_{ret_model} show ripple-like structure due to significant multipath. To model the LoS AoA using a given refractivity profile, the ray is traced using the interferometrically-measured observed AoA (as shown in Fig. 3). Since the
 320 observed AoA is noisy, the modelled LoS AoA will share similar noisy structure. The impact of multipath on the retrieved refractivity profiles is explored below. However, despite the presence of noise in the observed AoA measurement, the AoA_{ret_model} generally converges towards the AoA_{LoS} . The corresponding refractivity profiles for each 15-minute sample of observations between 08:36 and 09:36 UTC are shown in Fig. 10. The reference refractivity profile was determined using the 12:00 UTC Watnall radiosonde launch on the 22 September 2023. The refractivity value at the altitude of the receiver was fixed to the
 325 radiosonde value at an altitude of 0.575 km. The refractivity profiles retrieved using observational data shows agreement with the reference profile radiosonde profile in the lower atmosphere up to an altitude of ~ 5 km. At higher altitudes, the refractivity profiles show significant variability compared to lower altitudes, likely due to a lack of constraints here (a change in refractivity at higher altitudes has less impact than a change in refractivity at lower altitudes). The potential systematic uncertainties in the observations suggested by these results are explored in Section 4.

330 The difference between the AoA_{LoS} and AoA_{init_model} and the difference between the AoA_{LoS} and AoA_{ret_model} as a function of surface distance for the four observational periods is shown in Fig. 11. The distribution in the difference between the AoA_{LoS} and AoA_{init_model} and the difference between the AoA_{LoS} and AoA_{ret_model} is shown in Fig. 12. The mean of the LoS AoA difference decreased from $\sim 0.034^\circ$ to $\sim 0.009^\circ$ and the standard deviation increased from $\sim 0.023^\circ$ to $\sim 0.024^\circ$ before and after the retrieval respectively. The standard deviation in the LoS AoA difference exceeds the target accuracy of 0.01° ,
 335 meaning that complex refractivity structure is unlikely to be resolved. The AoA_{init_model} is larger than the AoA_{LoS} , indicating the initial atmosphere is less refractive than the true atmosphere. The retrieved atmosphere results in apparent weaker refraction of ADS-B signals from nearby aircraft and apparent stronger refraction for signals originating from more distant aircraft. The apparent variability in the observed refraction of ADS-B signals originating from aircraft at varying distances is likely to be partially a consequence of observational noise. The noise is visible in Fig. 9 as a ripple-like structure in the flight paths. This
 340 noise is correlated with the foreground environment and is not truly random, which likely results in a bias in the retrieved refractivity profiles. Potential sources of uncertainty, including multipath interference and instrumental errors are described in Section 4.

4 Discussion

Initial tests with synthetic data and the case study using actual measurements have revealed possible weaknesses in the retrieval
 345 method. Sources of experimental uncertainty that require further attention are discussed in the subsections below.

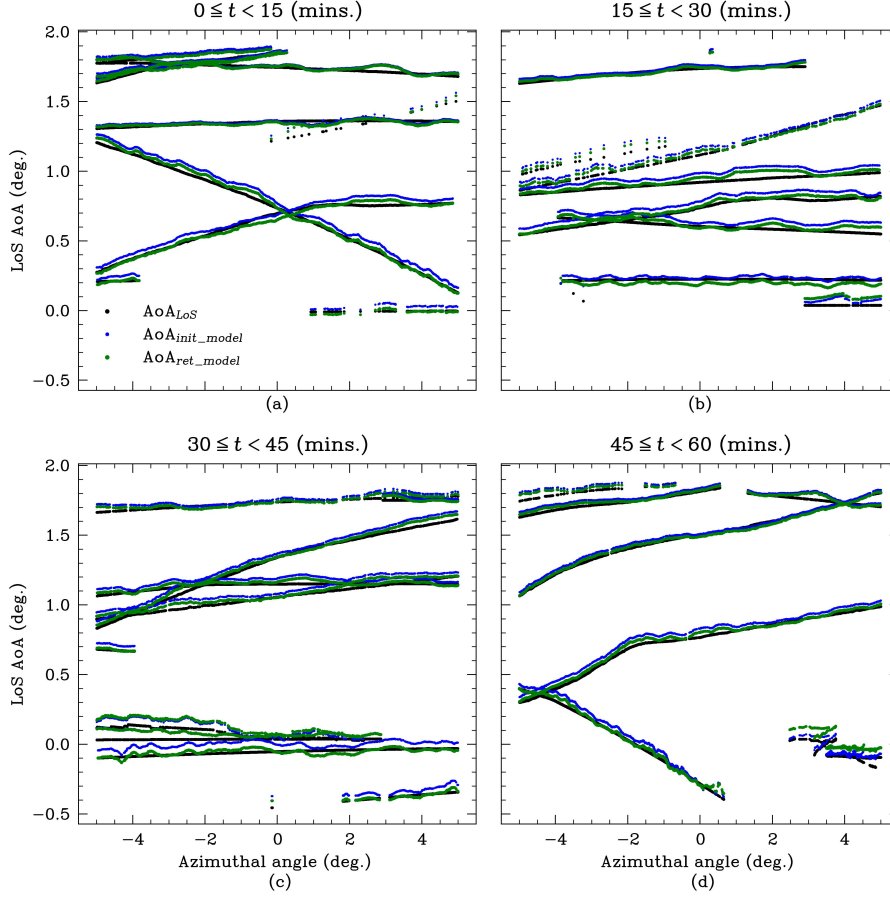


Figure 9. The LoS AoA as a function of azimuthal angle (where 0° corresponds to the centre of the observational sector). The black points (AoA_{LoS}) indicate the LoS AoA derived from the decoded ADS-B transmissions. The blue points indicate the modelled LoS AoA (AoA_{init_model}) determined from tracing rays through the initial (first-guess) refractivity profile. The green points indicate the modelled LoS AoA (AoA_{ret_model}) determined from tracing rays through the retrieved (optimised) refractivity profile. The time period was between 08:36 and 09:36 UTC.

4.1 Assumption of a horizontally homogeneous refractivity field

The SODE ray tracing model was derived with the assumption of spherical symmetry (the refractivity varies only with radial distance and is horizontally homogeneous). Therefore only one-dimensional refractivity profiles can be derived using the adjoint model presented here. For rays with negative elevation angles, this assumption will likely break down rapidly as the ray propagates through the highly variable planetary boundary layer. In this study only ADS-B transmissions with an observed $AoA \geq 0^\circ$ were used. As the rays generally increase in altitude with distance due to the curvature of the Earth, the assumption of spherical symmetry is therefore assumed to be reasonable. However, significant horizontal variability in refractivity can occur




Figure 10. The 12:00 UTC 22 September 2023 Watnall radiosonde refractivity profile (dotted), retrieved (solid colour) and initial (dashed) refractivity profiles. The time period was between 08:36 and 09:36 UTC. The four observational periods contained 9700 (08:36-08:51 UTC), 8315 (08:51-09:06 UTC), 8185 (09:06-09:21 UTC) and 7589 (09:21-09:36 UTC) received ADS-B transmissions respectively within the sector.

along frontal systems, which can invalidate the spherical symmetry assumption. An occluded front moved southeastwards during the observation period, which may have introduced some stronger horizontal variability than usual. The weather during
355 the observation period was partially cloudy, with patchy low cloud present around Clee Hill. Local variations in the refractivity near the receiver could contribute to uncertainty in the refractivity retrievals. However, the significant multipath contamination of the observed AoA measurement was expected to be the dominant source of uncertainty in the retrieved refractivity profiles.

To assimilate ADS-B refraction measurements influenced by significant horizontal refractivity gradients, a two-dimensional ray tracing model and its adjoint will be required. It is uncertain how well constrained a two-dimensional refractivity retrieval
360 will be using a single ADS-B interferometer. A network of ADS-B interferometers across the UK could potentially allow for two- or three-dimensional refractivity profiles to be obtained.

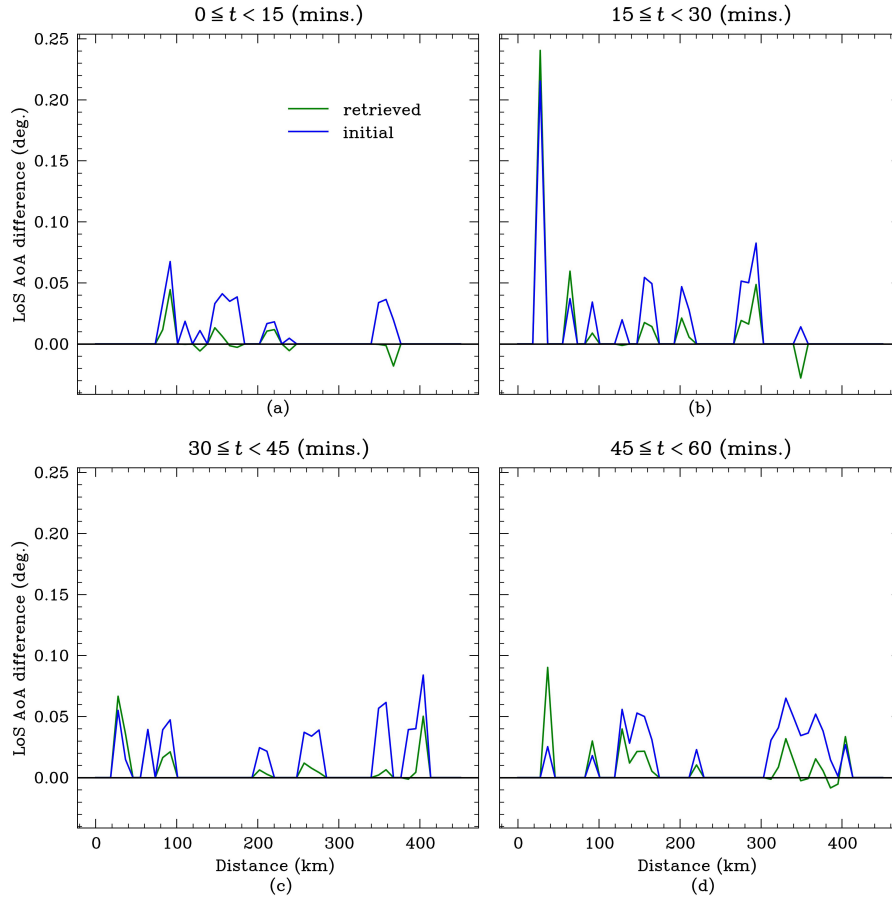


Figure 11. The difference between the AoA_{LoS} and AoA_{init_model} (blue line) and the difference between the AoA_{LoS} and AoA_{ret_model} (green line) as a function of distance. The time period was between 08:36 and 09:36 UTC.

4.2 Sensitivity of the retrievals to surface refractivity

The retrieved refractivity profile is extremely sensitive to the surface refractivity (as shown in Fig. 5). Uncertainty in the measured refractivity here will result in erroneous retrievals, therefore an accurate in-situ measurement of refractivity will likely be required to constrain the retrieved solution. Accurate on-site refractivity measurements could be obtained using a permanent weather station. The bending of radio signals is dominated by the refractivity gradient, rather than the absolute refractivity (since the refractive index of air is approximately unity even near the surface). Any uncertainty in the surface refractivity introduces a constant offset throughout the entire retrieved profile, however the retrieved refractivity gradient remains unchanged. The surface refractivity uncertainty could be addressed by instead assimilating refractivity gradients into NWP models, which would reduce the need for accurate surface refractivity measurements.




Figure 12. The distribution in the modelled LoS AoA difference for all observations using the initial (blue) and retrieved (green) refractivity profiles. The mean difference in the LoS AoA decreased from $\sim 0.034^\circ$ to $\sim 0.009^\circ$ for the initial and retrieved refractivity profiles respectively. The standard deviation in the LoS AoA difference increased from $\sim 0.023^\circ$ to $\sim 0.024^\circ$ for the initial and retrieved refractivity profiles respectively.

4.3 Multipath

Estimating the observed AoA to an uncertainty of 0.01° requires accurate measurements of the phase difference across the interferometer elements. A significant source of uncertainty in this measurement is due to the presence of multipath. At grazing elevation (observed AoA $\lesssim 2.0^\circ$) the direct and reflected ADS-B signals are received in rapid succession, resulting in mutual interference. The multipath results in ripple-like structures in the observed flight paths of the aircraft, as shown in Fig. 9. The uncertainties in the observed AoA measurements are highly correlated and dependent on the local terrain, which makes estimating the observation error covariance matrix very difficult. Efforts to minimise the impacts of multipath have included beam forming by stacking multiple antennas on each element of the interferometer (see Fig. 6) and increasing the sampling rate. The beam forming technique aims to deepen and control the position of the antenna nulls to minimise the sensitivity to foreground reflections. Increasing the sampling rate allows for phase measurements to be extracted from the direct signal before the arrival of the reflected signal. The sampling rate used for the experiment was 10 MHz, which was able to strongly suppress multipath originating from foreground reflections with a delay longer than ~ 100 ns (excess path length ~ 30 m). Multipath due to foreground reflections with a path excess shorter than this can only be partially suppressed using the current ADS-B interferometer. Future experiments will aim to further suppress the multipath by increasing sampling rate and modifying the beam pattern of the stacks. Both techniques have resulted in significant improvements in the accuracy of the observed AoA measurement. The multipath due to foreground reflections could also be mitigated using a long time series of observations. The impact of reflections from a static foreground would be a systematic perturbation to the observed phase difference. With enough observations from aircraft traversing at a range of elevations and distances the constant multipath interference could be

well characterised after accounting for atmospheric variability. The non-constant variability, such as atmospheric refractivity changes, could be de-trended out with sufficient observations. A similar technique is used in ground-based GNSS radio occultation, where Padullés et al. (2016) describes how time series of the observations can be used to remove a significant contribution of the multipath effect on the uncertainty in the measurements of the atmospheric polarimetric effects on observables. Future experiments using the prototype ADS-B interferometer will aim to quantify the nature of the observational noise arising due to multipath contamination.

4.4 Instrumental uncertainties

Instrumental errors may include the assumption that a measured phase difference of zero corresponds to an observed AoA of 0.00° . A slight offset between the upper and lower elements of the interferometer could introduce uncertainty in this assumption. The reported position of the aircraft is also limited by the reporting resolution (on the order ~ 10 m, Stone and Kitchen (2015)), however the impact of positional uncertainty on the observed AoA was expected to be minor due to the large distances to aircraft (Lewis et al., 2023a). Uncertainty in the altitude of the receiver will introduce a systematic error in the retrieved refractivity profiles. Figure 8 shows the impact of uncertainty in the height of the receiver on the retrieved refractivity profiles. The height uncertainty results in significant variability in refractivity at high altitudes. The weakly constrained refractivity retrievals above an altitude of ~ 5 km for some of the retrieved profiles in Fig. 10 suggest significant observational noise. This suggests a potential systematic error in the observational data. However, since the high altitude variability was only present in some of the retrieved profiles, it may suggest that the source of the uncertainty may not be entirely systematic. The gradient of the penalty function with respect to refractivity at higher altitudes is significantly lower than at lower altitudes, as shown in Fig. 5. Introducing additional constraints to the penalty function to penalise large departures from the reference refractivity profile may reduce the presence of unphysical refractivity structures. In practice, a one-dimensional variational (1D-Var) data assimilation scheme would prevent unphysical retrievals. However, a good understanding of the observational uncertainties would be required.

The mean altitude of the midpoint between the elements of the interferometer above the ellipsoid Earth model was estimated to be 575 m using a GNSS receiver located on site. The standard deviation of the measured altitude was estimated to be 2 m. Multipath can result in correlated uncertainties in the observed AoA measurements, potentially resulting in biased retrievals.

5 Conclusions

A new method has been developed to retrieve refractivity profiles from the interferometrically-measured AoA of radio signals from aircraft. The method incorporates the SODE ray tracing model (Zeng et al., 2014) and was tested on both synthesised data (to explore sensitivity to measurement noise) and real observational data measured with an experimental interferometer. Three AoA measurement noise cases were simulated, consisting of no noise and two normally distributed noise cases (with standard deviations of 0.01° and 0.05° respectively) injected into the observed AoA measurements. The impact of each of the injected observed AoA measurement uncertainty cases on the retrieved vertical refractivity profiles was investigated. The convergence

of the penalty function was impacted negatively and the RMSE in the retrieved refractivity profile increased as the standard deviation in the injected (normally distributed) noise increased. The synthetic retrievals using an AoA noise standard deviation of 0.01° were of comparable accuracy to vertical refractivity profile retrievals using GNSS radio occultation sounding (Anthes et al., 2022), airborne GNSS radio occultation sounding (Murphy et al., 2015) and GNSS troposphere tomography (Wu et al., 2014; Trzcina and Rohm, 2019). When the algorithm was tested on real measurements in four separate 15-minute periods, the retrieved profiles showed significant variability at high altitudes, but good convergence towards the reference radiosonde profile in the lowest few kilometres above the surface. Sources of uncertainty were explored and future work was suggested, such as improvements in the instrument design to minimise multipath contamination and the development of a two-dimensional adjoint model to capture horizontal variations in refractivity.

430 Appendix A

A more general constrained optimisation problem can be expressed in terms of the state variables \mathbf{v} and the control variables \mathbf{p} as

$$\min_{\mathbf{p}} \quad f(\mathbf{v}(\mathbf{p})) \quad (\text{A1a})$$

$$\text{subject to} \quad \mathbf{g}(\mathbf{v}, \mathbf{p}) = 0, \quad (\text{A1b})$$

435

where f is the scalar functional we seek to minimise and \mathbf{g} are the constraints on the optimisation (equivalent to the constraint equations (11b)-(11e) in the original optimisation problem) respectively. In a data assimilation approach, \mathbf{g} is sometimes referred to as the forward model and describes the physics of the system. We can convert the constrained optimisation problem described by (A1) into an unconstrained optimisation problem by introducing two vectors of slack variables known as Lagrange multipliers. The Lagrangian, \mathcal{L} , can be defined in terms of f and the constraints \mathbf{g} as

440

$$\mathcal{L}(\mathbf{v}, \mathbf{p}, \mathbf{\Lambda}) \equiv f(\mathbf{v}) - \langle \mathbf{\Lambda}, \mathbf{g}(\mathbf{v}, \mathbf{p}) \rangle, \quad (\text{A2})$$

where $\mathbf{\Lambda}$ is a vector of adjoint variables associated with the constraints \mathbf{g} and $\langle \cdot, \cdot \rangle$ is the inner product. The minimum of the Lagrangian, \mathcal{L} , and the scalar functional, f , will be equivalent when all the variables except for \mathbf{p} are stationary points (that is, variations of the Lagrangian with respect to \mathbf{v} and $\mathbf{\Lambda}$ are equal to zero). The total derivative of \mathcal{L} with respect to \mathbf{p} is written as

$$\frac{d\mathcal{L}}{d\mathbf{p}} = \frac{\partial f}{\partial \mathbf{v}} \frac{d\mathbf{v}}{d\mathbf{p}} - \frac{d\mathbf{\Lambda}^T}{d\mathbf{p}} \mathbf{g} - \mathbf{\Lambda}^T \left(\frac{\partial \mathbf{g}}{\partial \mathbf{v}} \frac{d\mathbf{v}}{d\mathbf{p}} + \frac{\partial \mathbf{g}}{\partial \mathbf{p}} \right). \quad (\text{A3})$$

Since the constraints $\mathbf{g} = 0$, the second term on the right hand side of (A3) can be dropped. The expression then becomes

$$\frac{d\mathcal{L}}{d\mathbf{p}} = \frac{d\mathbf{v}}{d\mathbf{p}} \left(\frac{\partial f}{\partial \mathbf{v}} - \mathbf{\Lambda}^T \frac{\partial \mathbf{g}}{\partial \mathbf{v}} \right) - \mathbf{\Lambda}^T \frac{\partial \mathbf{g}}{\partial \mathbf{p}}. \quad (\text{A4})$$

In a data assimilation scheme, the Jacobian matrix $d\mathbf{v}/d\mathbf{p}$ describes the sensitivity of the trajectory of the state variables \mathbf{v} in the forward model to changes in the control variables \mathbf{p} . For large numbers of control variables the cost of computing $d\mathbf{v}/d\mathbf{p}$ quickly becomes very high. Defining the adjoint variable $\mathbf{\Lambda}$ such that

$$\mathbf{\Lambda} \left(\frac{\partial \mathbf{g}}{\partial \mathbf{v}} \right)^T = \left(\frac{\partial f}{\partial \mathbf{v}} \right)^T, \quad (\text{A5})$$

the first term on the right hand side of (A4) is zero and the need to calculate the Jacobian matrix $d\mathbf{v}/d\mathbf{p}$ is avoided (Zou et al., 1997; Gill et al., 1981).

The penalty function for a single ray, j , can be augmented with the known dynamics of the system using Lagrange multipliers

$$\mathcal{L} = j - \int_0^{r_f} \psi (\dot{h} - u) dr - \int_0^{r_f} \mu \left(\dot{u} - (1 - u^2) \left(\frac{1}{n} \frac{dn}{dh} + \frac{1}{a + h} \right) \right) dr, \quad (\text{A6})$$

where $\dot{\psi}$ and $\dot{\mu}$ are the derivatives of the adjoint variables ψ and μ with respect to r respectively. The stationary points of the Lagrangian with respect to h and u are found as

$$\delta \mathcal{L}|_h = \delta j|_h - \int_0^{r_f} \psi \delta \dot{h} dr - \int_0^{r_f} \mu (1 - u^2) \left(\left(\frac{1}{n} \frac{dn}{dh} \right)^2 - \frac{1}{n} \frac{d^2 n}{dh^2} + \frac{1}{(a + h)^2} \right) \delta h dr = 0, \quad (\text{A7})$$

$$\delta \mathcal{L}|_u = \delta j|_u + \int_0^{r_f} \psi \delta u dr - \int_0^{r_f} \mu \delta \dot{u} dr - \int_0^{r_f} 2\mu u \left(\frac{1}{n} \frac{dn}{dh} + \frac{1}{a + h} \right) \delta u dr = 0. \quad (\text{A8})$$

Integrating (A7) and (A8) by parts

$$\delta \mathcal{L}|_h = \delta j|_h - [\psi \delta h]_0^{r_f} + \int_0^{r_f} \dot{\psi} \delta h dr - \int_0^{r_f} \mu (1 - u^2) \left(\left(\frac{1}{n} \frac{dn}{dh} \right)^2 - \frac{1}{n} \frac{d^2 n}{dh^2} + \frac{1}{(a + h)^2} \right) \delta h dr = 0, \quad (\text{A9})$$

$$\delta \mathcal{L}|_u = \delta j|_u + [\mu \delta u]_0^{r_f} + \int_0^{r_f} \dot{\mu} \delta u dr - \int_0^{r_f} 2\mu u \left(\frac{1}{n} \frac{dn}{dh} + \frac{1}{a + h} \right) \delta u dr = 0. \quad (\text{A10})$$

The adjoint variables at $r = 0$ are $\psi_{r=0} = 0$ and $\mu_{r=0} = 0$. Therefore

$$\psi|_{r=r_f} = \frac{\partial j}{\partial h}, \quad (\text{A11})$$

$$465 \quad \mu|_{r=r_f} = \frac{\partial j}{\partial u}, \quad (\text{A12})$$

$$\frac{d\psi}{dr} = \mu(1-u^2) \left(\left(\frac{1}{n} \frac{dn}{dh} \right)^2 - \frac{1}{n} \frac{d^2 n}{dh^2} + \frac{1}{(a+h)^2} \right), \quad (\text{A13})$$

$$\frac{d\mu}{dr} = 2\mu u \left(\frac{1}{n} \frac{dn}{dh} + \frac{1}{a+h} \right) - \psi. \quad (\text{A14})$$

Appendix B

The reported LoS AoA of the aircraft at the location of the observer is computed in terms of the geodetic Cartesian coordinates of the observer. The geodetic Cartesian coordinates of the observer (X_o, Y_o, Z_o) can be calculated using the geodetic coordinates latitude (ϕ_o), longitude (λ_o) and height (h_o) of the observer as

$$X_o = (N(\phi_o) + h_o) \cos(\phi_o) \cos(\lambda_o), \quad (\text{B1})$$

$$Y_o = (N(\phi_o) + h_o) \cos(\phi_o) \sin(\lambda_o), \quad (\text{B2})$$

475

$$Z_o = (N(\phi_o) + h_o) \sin(\phi_o), \quad (\text{B3})$$

where N is the prime vertical radius, defined in terms of the geodetic latitude ϕ as

$$N(\phi) = \frac{a}{\sqrt{1 - e^2 \sin^2(\phi)}}, \quad (\text{B4})$$

and $e^2 = 1 - b^2/a^2$, where $a = 6378.137$ km and $b = 6356.75231425$ km respectively (Decker, 1986; Hofmann-Wellenhof et al., 1997). The ellipsoid Earth geometry is shown in Fig. B1.

480

The geocentric coordinates of the aircraft (X'_p, Y'_p, Z'_p) are given by




Figure B1. The ellipsoid Earth model (thick grey line). The sphere with the radius equal to the prime vertical radius is shown.

$$X'_p = (N(\phi_p) + h_p)\cos(\phi_p)\cos(\lambda_p), \quad (\text{B5})$$

$$Y'_p = (N(\phi_p) + h_p)\cos(\phi_p)\sin(\lambda_p), \quad (\text{B6})$$

485

$$Z'_p = (N(\phi_p)(1 - e^2) + h_p)\sin(\phi_o), \quad (\text{B7})$$

where ϕ_p , λ_p and h_p are the geodetic latitude, geodetic longitude and altitude of the aircraft above the ellipsoid respectively.

The coordinates of the aircraft in terms of the geodetic coordinates of the observer are

$$X_p = X'_p, \quad (\text{B8})$$

490

$$Y_p = Y'_p, \quad (\text{B9})$$

$$Z_p = Z'_p + e^2 N(\phi_o)\sin(\phi_o). \quad (\text{B10})$$

The geometry is shown in Fig. B2. Defining the position vectors of the observer and aircraft as $\mathbf{r}_o = (X_o, Y_o, Z_o)$ and $\mathbf{r}_p =$
 495 (X_p, Y_p, Z_p) respectively, the vector between the position of the observer and the aircraft is

$$\Delta \mathbf{r} = \mathbf{r}_p - \mathbf{r}_o. \quad (\text{B11})$$

Therefore the zenith angle, Ω , of the aircraft as seen by the observer can be written in terms of the unit vectors

$$\Omega = \cos^{-1}(\Delta \hat{\mathbf{r}} \cdot \hat{\mathbf{r}}_o). \quad (\text{B12})$$

The reported LoS AoA of the aircraft is therefore equal to $\pi/2 - \Omega$. The mean local radius of curvature of an ellipsoid at a
 500 given point and azimuth direction is given by (Jekeli, 2006)

$$R_{\text{roc}} = \left(\frac{\sin^2(\alpha)}{N} + \frac{\cos^2(\alpha)}{M} \right)^{-1}, \quad (\text{B13})$$

where M is the radius of curvature of the meridian ellipse, given by

$$M = \frac{a(1 - e^2)}{(1 - e^2 \sin^2(\phi_o))^{\frac{3}{2}}}. \quad (\text{B14})$$

The radius of spherical Earth used in the ray tracing model was assumed to be the local radius of curvature of the ellipsoid at
 505 the position of the receiver, $R_{\text{roc}} \approx 6383.57$ km (assuming an azimuth $\alpha = 45^\circ$ north east).

Code and data availability. The ray tracing and adjoint inversion code are provided in the Github repository available at https://github.com/osl202/ADS_B_refractivity. The ADS-B data used in this study are available at <https://doi.org/10.5281/zenodo.12783749>. The radiosonde data was provided by the University of Wyoming (see <http://weather.uwyo.edu/upperair/sounding.html>).

Author contributions. Conceptualisation and instrument design, CB and MK; adjoint model design, OL; software, OL and CB; formal
 510 analysis, OL and CB; investigation, OL; data curation, CB; writing—original draft preparation, OL; writing—review and editing, MK, CB, NEB, EKS; visualisation, OL; supervision, CB, MK, NEB, EKS. All authors have read and agreed to the published version of the manuscript.

Competing interests. The contact author has declared that none of the authors have any competing interests.

Financial support. This work was supported by the Met Office, the University of Exeter and the Harry Otten Foundation.




Figure B2. The Z -coordinate of the aircraft in terms of the geodetic coordinates of the observer.

Acknowledgements. The authors would like to thank Met Office engineers Mike Prots and Paul Eadie for their support with the installation of the interferometer on the Clee Hill weather radar tower. The authors would also like to thank Prof Matthew Browning (University of Exeter) for the very helpful discussion on the adjoint method and Dr Claire Witham (Met Office) for offering feedback that significantly improved this manuscript. The authors would also like to thank the three anonymous referees for their insight and useful comments that have significantly improved the clarity of the manuscript. This project has made use of the following Python packages: PANDAS (McKinney, 2010), NUMPY (Harris et al., 2020), MATPLOTLIB (Hunter, 2007), SMPLOTLIB (Li, 2023), CARTOPY (Met Office, 2010 - 2015) and IRIS (Iris contributors).

References

- Anthes, R., Sjoberg, J., Feng, X., and Syndergaard, S.: Comparison of COSMIC and COSMIC-2 radio occultation refractivity and bending angle uncertainties in August 2006 and 2021, *Atmosphere*, 13, 790, 2022.
- Aparicio, J. M. and Laroche, S.: An evaluation of the expression of the atmospheric refractivity for GPS signals, *Journal of Geophysical Research: Atmospheres*, 116, 2011.
- de Haan, S.: High-resolution wind and temperature observations from aircraft tracked by Mode-S air traffic control radar, *Journal of Geophysical Research: Atmospheres*, 116, 2011.
- de Haan, S. and Stoffelen, A.: Assimilation of high-resolution Mode-S wind and temperature observations in a regional NWP model for nowcasting applications, *Weather and Forecasting*, 27, 918–937, 2012.
- de Haan, S., de Haij, M., and Soudij, J.: The use of a commercial ADS-B receiver to derive upper air wind and temperature observations from Mode-S EHS information in The Netherlands, KNMI De Bilt, The Netherlands, 2013.
- Decker, B. L.: World geodetic system 1984, Defense Mapping Agency Aerospace Center St Louis Afs Mo, 1986.
- Fabry, F., Frush, C., Zawadzki, I., and Kilambi, A.: On the extraction of near-surface index of refraction using radar phase measurements from ground targets, *Journal of Atmospheric and Oceanic Technology*, 14, 978–987, 1997.
- Gaffard, C., Li, Z., Harrison, D., Lehtinen, R., and Roininen, R.: Evaluation of a Prototype Broadband Water-Vapour Profiling Differential Absorption Lidar at Cardington, UK, *Atmosphere*, 12, 1521, 2021.
- Geer, A. J., Baordo, F., Bormann, N., Chambon, P., English, S. J., Kazumori, M., Lawrence, H., Lean, P., Lonitz, K., and Lupu, C.: The growing impact of satellite observations sensitive to humidity, cloud and precipitation, *Quarterly Journal of the Royal Meteorological Society*, 143, 3189–3206, 2017.
- Gill, P. E., Murray, W., and Wright, M. H.: Practical optimization, Academic Press Inc. [Harcourt Brace Jovanovich Publishers], London, ISBN 0-12-283950-1, 1981.
- Guo, X., Wu, D., Wang, Z., Wang, B., Li, C., Deng, Q., and Liu, D.: A review of atmospheric water vapor lidar calibration methods, *WIREs Water*, n/a, e1712, <https://doi.org/https://doi.org/10.1002/wat2.1712>, 2023.
- Harris, C. R., J., M. K., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., H., v. M., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E.: Array programming with NumPy, *Nature*, 585, 357–362, <https://doi.org/10.1038/s41586-020-2649-2>, 2020.
- Healy, S. B. and Thépaut, J.-N.: Assimilation experiments with CHAMP GPS radio occultation measurements, *Quarterly Journal of the Royal Meteorological Society*, 132, 605–623, 2006.
- Held, I. M. and Soden, B. J.: Water vapor feedback and global warming, *Annual review of energy and the environment*, 25, 441–475, 2000.
- Hofmann-Wellenhof, B., Lichtenegger, H., and Collins, J.: Global Positioning System: Theory and Practice, Springer-Verlag, ISBN 9783211828397, <https://books.google.co.uk/books?id=iZTAAAAMAAJ>, 1997.
- Hunter, J. D.: Matplotlib: A 2D graphics environment, *Computing in Science & Engineering*, 9, 90–95, <https://doi.org/10.1109/MCSE.2007.55>, 2007.
- ICAO: Technical Provisions for Mode S Services and Extended Squitter. 2nd ed. International Civil Aviation Organization Doc. 9871, AN/460, 2012.

- Ingleby, B., Isaksen, L., and Kral, T.: Evaluation and impact of aircraft humidity data in ECMWF's NWP system, European Centre for Medium-Range Weather Forecasts, 2019.
- Iris contributors: Iris, <https://doi.org/10.5281/zenodo.595182>.
- 560 Jekeli, C.: Geometric reference systems in geodesy, Division of Geodesy and Geospatial Science, School of Earth Sciences, Ohio State University, 25, 2006.
- Kingma, D. P. and Ba, J.: Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980, 2014.
- Kirkpatrick, S., Gelatt Jr, C. D., and Vecchi, M. P.: Optimization by simulated annealing, *Science*, 220, 671–680, 1983.
- Kuo, Y.-H., Sokolovskiy, S. V., Anthes, R. A., and Vandenberghe, F.: Assimilation of GPS radio occultation data for numerical weather
565 prediction, *Terrestrial Atmospheric and Oceanic Sciences*, 11, 157–186, 2000.
- Kursinski, E. R., Hajj, G. A., Schofield, J. T., Linfield, R. P., and Hardy, K. R.: Observing Earth's atmosphere with radio occultation measurements using the Global Positioning System, *Journal of Geophysical Research: Atmospheres*, 102, 23 429–23 465, 1997.
- Lewis, O., Brunt, C., and Kitchen, M.: A new method of retrieving atmospheric refractivity structure, *International Journal of Remote Sensing*, 44, 749–785, 2023a.
- 570 Lewis, O., Brunt, C., Kitchen, M., Healy, S., and Bowler, N.: A new method of retrieving atmospheric refractivity, in: *Remote Sensing of Clouds and the Atmosphere XXVIII*, vol. 12730, pp. 137–145, SPIE, 2023b.
- Li, J.: smplotlib, <https://doi.org/10.5281/zenodo.8126529>, 2023.
- Li, Z.: Impact of assimilating Mode-S EHS winds in the Met Office's high-resolution NWP model, *Meteorological Applications*, 28, e1989, 2021.
- 575 McKinney, W.: Data Structures for Statistical Computing in Python, in: *Proceedings of the 9th Python in Science Conference*, edited by Stéfan van der Walt and Jarrod Millman, pp. 56 – 61, <https://doi.org/10.25080/Majora-92bf1922-00a>, 2010.
- Met Office: Cartopy: a cartographic python library with a Matplotlib interface, <https://scitools.org.uk/cartopy>, 2010 - 2015.
- Mirza, A. K., Ballard, S. P., Dance, S. L., Maisey, P., Rooney, G. G., and Stone, E. K.: Comparison of aircraft-derived observations with in situ research aircraft measurements, *Quarterly Journal of the Royal Meteorological Society*, 142, 2949–2967, 2016.
- 580 Mirza, A. K., Ballard, S. P., Dance, S. L., Rooney, G. G., and Stone, E. K.: Towards operational use of aircraft-derived observations: a case study at London Heathrow airport, *Meteorological Applications*, 26, 542–555, 2019.
- Möller, G.: Reconstruction of 3D wet refractivity fields in the lower atmosphere along bended GNSS signal paths, Doctoral dissertation, Wien, 2017.
- Murphy, B. J., Haase, J. S., Muradyan, P., Garrison, J. L., and Wang, K.-N.: Airborne GPS radio occultation refractivity profiles observed in
585 tropical storm environments, *Journal of Geophysical Research: Atmospheres*, 120, 1690–1709, 2015.
- Padullés, R., Cardellach, E., de la Torre Juárez, M., Tomás, S., Turk, F. J., Oliveras, S., Ao, C. O., and Rius, A.: Atmospheric polarimetric effects on GNSS radio occultations: The ROHP-PAZ field campaign, *Atmospheric Chemistry and Physics*, 16, 635–649, 2016.
- Pauley, P. M. and Ingleby, B.: Assimilation of in-situ observations, *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications* (Vol. IV), pp. 293–371, 2022.
- 590 Petersen, R. A., Counce, L., Mamrosh, R., Baker, R., and Pauley, P.: On the impact and future benefits of AMDAR observations in operational forecasting: Part II: Water vapor observations, *Bulletin of the American Meteorological Society*, 97, 2117–2133, 2016.
- Rieckh, T., Anthes, R., Randel, W., Ho, S.-P., and Foelsche, U.: Evaluating tropospheric humidity from GPS radio occultation, radiosonde, and AIRS from high-resolution time series, *Atmospheric Measurement Techniques*, 11, 3091–3109, 2018.
- Rohm, W.: The precision of humidity in GNSS tomography, *Atmospheric research*, 107, 69–75, 2012.

- 595 Smith, E. K. and Weintraub, S.: The constants in the equation for atmospheric refractive index at radio frequencies, *Proceedings of the IRE*, 41, 1035–1037, 1953.
- Stone, E. K.: A comparison of Mode-S Enhanced Surveillance observations with other in situ aircraft observations, *Quarterly Journal of the Royal Meteorological Society*, 144, 695–700, 2018.
- Stone, E. K. and Kitchen, M.: Introducing an approach for extracting temperature from aircraft GNSS and pressure altitude reports in ADS-B
600 messages, *Journal of Atmospheric and Oceanic Technology*, 32, 736–743, 2015.
- Stone, E. K. and Pearce, G.: A network of Mode-S receivers for routine acquisition of aircraft-derived meteorological data, *Journal of Atmospheric and Oceanic Technology*, 33, 757–768, 2016.
- Tang, Y., Lean, H. W., and Bornemann, J.: The benefits of the Met Office variable resolution NWP model for forecasting convection, *Meteorological Applications*, 20, 417–426, 2013.
- 605 Teh, A., O’Toole, M., and Gkioulekas, I.: Adjoint nonlinear ray tracing, *ACM Transactions on Graphics (TOG)*, 41, 1–13, 2022.
- Trenberth, K. E., Fasullo, J., and Smith, L.: Trends and variability in column-integrated atmospheric water vapor, *Climate Dynamics*, 24, 741–758, 2005.
- Trenberth, K. E., Fasullo, J. T., and Kiehl, J.: Earth’s global energy budget, *Bulletin of the American Meteorological Society*, 90, 311–324, 2009.
- 610 Trzcina, E. and Rohm, W.: Estimation of 3D wet refractivity by tomography, combining GNSS and NWP data: First results from assimilation of wet refractivity into NWP, *Quarterly Journal of the Royal Meteorological Society*, 145, 1034–1051, 2019.
- Wang, K.-N., Ao, C. O., Morris, M. G., Hajj, G. A., Kurowski, M. J., Turk, F. J., and Moore, A. W.: Joint 1DVar retrievals of tropospheric temperature and water vapor from Global Navigation Satellite System radio occultation (GNSS-RO) and microwave radiometer observations, *Atmospheric Measurement Techniques*, 17, 583–599, 2024.
- 615 Weckwerth, T. M., Pettet, C. R., Fabry, F., Park, S. J., LeMone, M. A., and Wilson, J. W.: Radar refractivity retrieval: Validation and application to short-term forecasting, *Journal of Applied Meteorology and Climatology*, 44, 285–300, 2005.
- Wu, X., Wang, X., and Lü, D.: Retrieval of vertical distribution of tropospheric refractivity through ground-based GPS observation, *Advances in Atmospheric Sciences*, 31, 37–47, 2014.
- Xie, F., Syndergaard, S., Kursinski, E. R., and Herman, B. M.: An approach for retrieving marine boundary layer refractivity from GPS
620 occultation data in the presence of superrefraction, *Journal of Atmospheric and Oceanic Technology*, 23, 1629–1644, 2006.
- Zeng, Y., Blahak, U., Neuper, M., and Jerger, D.: Radar beam tracing methods based on atmospheric refractive index, *Journal of Atmospheric and Oceanic Technology*, 31, 2650–2670, 2014.
- Zhao, Q., Zhang, K., and Yao, W.: Influence of station density and multi-constellation GNSS observations on troposphere tomography, in: *Annales Geophysicae*, vol. 37, pp. 15–24, Copernicus Publications Göttingen, Germany, 2019.
- 625 Zou, X., Vandenbergh, F., Pondeva, M., and Kuo, Y.-H.: Introduction to adjoint techniques and the MM5 adjoint modeling system, NCAR Technical note, 1997.