

Statistics of travelling ionospheric disturbances observed using the LOFAR radio telescope

Ben Boyde^{1,*} , Alan G. Wood¹ , Gareth Dorrian¹ , Francesco de Gasperin^{2,3} , and Maaijke Mevius⁴ 

¹ School of Engineering, University of Birmingham, Edgbaston, B15 2TT Birmingham, UK

² INAF – Istituto di Radioastronomia, Via P. Gobetti 101, 40129 Bologna, Italy

³ Hamburger Sternwarte, Universität Hamburg, Gojenbergsweg 112, D-21029, Hamburg, Germany

⁴ ASTRON – The Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands

Received 28 August 2024 / Accepted 28 January 2025

Abstract—A climatology of Travelling Ionospheric Disturbances (TIDs) observed using the Low Frequency ARray (LOFAR) has been created based on 2723 h of astronomical observations. Radio telescopes such as LOFAR must contend with many causes of signal distortions, including the ionosphere. To produce accurate astronomical images, calibration solutions are derived to mitigate these distortions as much as possible. These calibration solutions provide extremely precise measurements of ionospheric variations across the LOFAR network, enabling TIDs to be detected which may be inaccessible to more traditional techniques. Waves are detected by LOFAR under all observing conditions, with no clear dependence on solar or geomagnetic activity. The vast majority of the observed waves travel in the opposite direction to the climatological thermospheric winds, suggesting that they are caused by upward propagating atmospheric gravity waves which are filtered by the wind. Waves of different periods display slightly different propagation directions, with waves of shorter periods consistent with the winds at lower altitudes within the thermosphere (180 km for 10–15 min periods compared to 220 km for 20–27 min periods). This suggests that either the shorter period waves are being detected at lower altitudes or that they are simply more sensitive to the winds at lower altitudes. This indicates that observations made using LOFAR may enable the investigation of vertical coupling within the neutral atmosphere. The shortest period waves in the dataset ($< \sim 10$) display distinct characteristics, suggesting they may be from a distinct population such as previously reported disturbances in the plasmasphere. The short period waves are compared to previous observations using other radio telescopes, showing that plasmaspheric disturbances likely account for some of the shortest period waves ($< \sim 5$ min) but there are still a large number of waves at these periods which are of uncertain origin.

Keywords: Travelling ionospheric disturbance / Atmospheric gravity wave / Thermosphere-ionosphere coupling

1 Introduction

Travelling ionospheric disturbances (TIDs) are wave-like disturbances in the ionosphere, with typical wavelengths ranging from 10 s to 1000 s of kilometres and typical velocities of around $50\text{--}500\text{ ms}^{-1}$ (Hocke & Schlegel, 1996). They have been observed with a wide range of instruments, including ionosondes (e.g. Negrea et al., 2016, 2018), incoherent scatter radars (e.g. Oliver et al., 1997; Galushko et al., 1998; Panasenko et al., 2018), HF radars (e.g. Frissell et al., 2014, 2016; Oinats et al., 2016), HF Doppler sounders (e.g. Waldock & Jones, 1986; Crowley et al., 1987; Chum et al., 2010; Fišer et al., 2017), satellite measurements (e.g. Yin et al., 2019), Global

Navigation Satellite Systems (GNSS) Total Electron Content (TEC) measurements (e.g. Ding et al., 2011; Otsuka et al., 2013; Sivakandan et al., 2021), radio interferometers (e.g. Jacobson et al., 1995), and radio telescopes (e.g. Koval et al., 2017; Fallows et al., 2020; Boyde et al., 2022; Dorrian et al., 2023). They are mainly caused by atmospheric gravity waves (AGWs) (Hines, 1960), perturbations of the neutral atmosphere that cause corresponding perturbations to the ionospheric plasma.

AGWs in the thermosphere have a range of sources, both within the thermosphere itself and lower in the atmosphere. Lower atmospheric sources include wind flow over orography (e.g. Liu et al., 2019), convective systems (e.g. Lane et al., 2001), large scale structures such as jet streams and the polar vortex (e.g. Buss et al., 2004; Sato & Yoshiki, 2008), as well

*Corresponding author: bxbl57@student.bham.ac.uk

as impulsive events such as tsunamis, earthquakes, and volcanoes (e.g. Laughman et al., 2017; Inchin et al., 2020; Wright et al., 2022). Generation in the thermosphere can occur due to Joule heating during geomagnetic storms (e.g. Hunsucker, 1982; Hocke & Schlegel, 1996), excitation of higher-order waves due to the dissipation of waves from lower in the atmosphere (e.g. Vadas et al., 2003), and non-linear wave-wave interactions (e.g. Wüst & Bittner, 2006). Upward propagating AGWs from the lower atmosphere transfer energy and momentum upwards, and are critical to understanding vertical coupling in the neutral atmosphere and the energy and momentum budget of the upper atmosphere (Yiğit et al., 2016). Even relatively small amplitude AGWs generated in the lower atmosphere can have significant effects in the thermosphere, as the amplitude increases exponentially with height to conserve the energy flux as the neutral density decreases. This increasing amplitude can eventually cause the wave to become unstable and break, depositing its energy and momentum which can then lead to excitation of higher-order AGWs (Vadas et al., 2003), as well as having significant influence on the circulation and thermal structure of the middle atmosphere (Fritts & Alexander, 2003). The neutral wind and density perturbations induced by the AGW cause perturbations in the electron density of the ionosphere, both by compression and rarefaction of existing plasma and changes in production and recombination rates, which are observed as TIDs (Hooke, 1968). Due to the difficulty of observing the neutral thermosphere directly, TIDs are often used as proxies to understand AGW propagation (e.g. Chum et al., 2021).

Upwards propagating AGWs are strongly influenced by the background atmospheric state. For example, the neutral winds apply a strong filtering effect which typically becomes significant around the mesosphere and lower thermosphere, such that AGWs propagating against the wind direction can propagate to higher altitudes and AGWs propagating with the wind cannot (Cowling et al., 1971). This directional preference has been reported in TID observations by several authors (e.g. Waldock & Jones, 1986; Crowley et al., 1987; Oinats et al., 2016; Zalozovski et al., 2021), although some have found that the dominant TID propagation direction can be systematically offset from the anti-wind direction by up to 90° while still tracking its diurnal evolution (e.g. Crowley & Rodrigues, 2012). As well as the neutral wind, variations in density, composition and temperature of the neutral atmosphere with altitude can lead to AGW dissipation due to kinematic viscosity or thermal diffusivity (e.g. Vadas, 2007). These effects are dependent on wave parameters such as period and vertical and horizontal wavelengths, meaning that they also limit the range of wave parameters we expect to observe in the thermosphere.

Many techniques for investigating TIDs and their relationship to AGWs are limited by not being able to resolve the full three-dimensional structure of the TID. However, some studies using incoherent scatter radar (e.g. Nicolls & Heinselman, 2007; Medvedev et al., 2015, 2017) and multi-frequency HF Doppler sounding (e.g. Chum et al., 2021) have succeeded in resolving the full three-dimensional wavevector. A comparison of these observations with theoretical TID dispersion relations, such as those developed by Hines (1960) and Vadas & Fritts (2005), enables direct verification of the causative relationship between TIDs and AGWs on a case study basis. Results from Irkutsk Incoherent Scatter Radar indicate that the majority of

TIDs are consistent with AGW driving, and the AGW related TIDs are particularly dominant during the daytime (Medvedev et al., 2017). These studies typically use statistical models for the atmospheric density and temperature profiles but allow the neutral wind to be a free parameter when fitting the AGW dispersion relation. The inferred neutral winds show general consistency with statistical models when taken as monthly medians (e.g. Medvedev et al., 2015, 2017) but significant variability in individual cases (e.g. Nicolls & Heinselman, 2007).

The LOw Frequency ARray (LOFAR: van Haarlem et al., 2013) is a radio telescope centred in the Netherlands capable of observing between 10–90 MHz (Low Band Array, LBA) and 110–250 MHz (High Band Array, HBA). It is composed of independent stations, divided into core stations (24 within a ~ 2 km radius in the North-East of the Netherlands), remote stations (14 stations in the Netherlands up to ~ 50 km from the core), and international stations in other countries across Europe. The ionosphere causes significant distortion to radio signals that pass through it at LOFAR's observing frequencies, and these distortions have been used to identify a small-scale TID (Boyde et al., 2022), substructure within a medium-scale TID (Dorrian et al., 2023), and to observe the morphology and evolution of small scale perturbations in sporadic-E layers (Trigg et al., 2024; Wood et al., 2024). For astronomical imaging, the network of stations is used as a radio interferometer, which requires all distortions to the measured signal including the ionosphere to be corrected as far as possible (de Gasperin et al., 2019). The ionospheric corrections for a given observing direction vary for each station in the LOFAR network as well as with time. These calibration solutions from astronomical imaging observations provide measurements of differential TEC (dTEC) across the LOFAR network, which are made with a precision of $< 1\text{mTECu}$ (Mevius et al., 2016).

These LOFAR calibration solutions have been used to investigate ionospheric perturbations by several authors. For example, Mevius et al. (2016) investigated the spectral slope and anisotropy of the irregularities in data from the HBA and showed that the results were consistent with a mixture of Kolmogorov turbulence and coherent wave-like structures, with a preference for alignment with the geomagnetic field. De Gasperin et al. (2018) investigated calibration solutions derived from LBA data and found results consistent with Mevius et al. (2016), as well as showing that the results were consistent with nearby GNSS TEC measurements. Beser et al. (2022) showed that wave-like irregularities in two case study observations had directions consistent with variations in the geomagnetic field observed at nearby observatories.

Recently, Boyde et al. (2024) developed a method for extracting wavelength, propagation direction and wave amplitude from the calibration solutions. This method uses a wavelet transform to estimate wave power and phase on each baseline as a function of time and period, and then fits the expected relationship of phase and power with baseline length and orientation for a simple plane wave. Noise in the measurements is quantified to provide estimates of the uncertainty of the estimated wave parameters. Using synthetic data, this was shown to be capable of identifying waves with amplitudes as low as a few mTECu and with wavelengths up to ~ 500 . For the full details of this method, the interested reader is referred to Boyde et al. (2024). This method provides the possibility of




Figure 1. The distribution of observing hours within the LBA survey as a function of Kp index (top left), F10.7 solar flux index (bottom left), and UT on the horizontal axis and month on the vertical (right). Local time is approximately UT + 0.5 h. These distributions exclude observations for which the calibration solutions were too unreliable for analysis, see text for details.

using LOFAR calibration solutions to characterise TIDs over Europe.

This paper is structured as follows: first, in [Section 2](#) the data used is introduced along with the processing used to remove erroneous data points and identify wave events and their parameters. After this, in [Section 3](#) the key results such as occurrence rates and observed wave characteristics are presented. This is followed in [Section 4](#) by a discussion of the results in relation to previous observations and possible physical explanations, and then the conclusions in [Section 5](#).

2 Data and processing

For this study, dTEC solutions calculated from the observations made as part of the LOFAR LBA Survey ([de Gasperin et al., 2021](#)) were used. This dataset comprises 2723 one-hour observations using one of three calibrator sources (3C 196, 3C 295, or 3C 380) based on which was closest in the sky to the target field for the observation. Almost all these observations were made during daytime as this has been found to reduce the severity of ionospheric scintillation observed with LOFAR (Fallows, private communication) and allow more accurate imaging ([de Gasperin et al., 2021](#)). This survey started in 2017, but the vast majority of observations (>90%) were obtained between June 2020 and March 2023.

[Figure 1](#) shows the distribution of observing time as a function of the Kp index, F10.7 index and UT/season. This clearly illustrates the fact that observations were made during day time ([Fig. 1](#), right panel), with good coverage across almost all seasons except for the lack of observations during May. The left panels show that observations were predominantly made

during low geomagnetic activity ([Fig. 1](#), top left) and low solar activity ([Fig. 1](#), bottom left). This was not due to deliberate scheduling during quiet conditions, but rather representative of the typical conditions during the period in which the observations were made, spanning the later part of solar minimum and the start of the rising phase of solar cycle 25.

2.1 Data flagging and interpolation

Before using these dTEC solutions to derive wave parameters, the data was first processed to remove errors which would distort the results. Firstly, for some baselines at some time steps, the calibration solution failed to find a dTEC value and these points were flagged within the data. Besides these flagged failed solution values, it was found that occasionally the dTEC value on a given baseline would change at a single time step by up to ~ 1 TECu before returning to its previous trend, an obviously unphysical variation. In order to identify these instances, the first difference of the dTEC series on each baseline was taken, and points in this series which had a magnitude exceeding 5 standard deviations were identified as potential spikes, analogous to the approach used to identify radio frequency interference in LOFAR scintillation data ([Fallows et al., 2020](#)). If two consecutive points in the first difference time series exceeded 5 standard deviations and had opposite signs, the corresponding point in the dTEC series was flagged as a bad data point. If a given baseline had more than 5% of its data points flagged, it was excluded from further analysis. Observations with more than 40% of baselines removed in this way were not considered in the analysis, which amounted to 94 observations being discarded leaving 2629 one-hour observations for the analysis. Almost all observations required some

data points to be replaced in this way, but the observations which were entirely removed were predominantly at low elevation ($< \sim 40^\circ$).

For baselines that were not removed from the analysis, the flagged data points were replaced by linear interpolation from neighbouring values, or extrapolation with a constant value if they occurred at either end of the time series. This replacement of bad data was necessary as the spectral analysis method used requires uniformly sampled time series of equal length on all baselines. Once the bad data points were replaced, the values were rescaled from differential slant TEC to the corresponding differential vertical TEC assuming a thin ionospheric shell at an altitude of 350 km. This correction is not exact for dTEC measurements due to the differences in elevation between different stations, but given that baselines are no more than ~ 50 km which corresponds to around a tenth of a degree, this is negligible for baselines within the Dutch LOFAR network.

2.2 Event identification

The method presented in Boyde et al. (2024) calculates the wave properties from the calibration solutions for a specified time and wave period. Given that these cannot be manually defined for analysing such a large dataset, it was necessary to define a sampling approach for selecting the times and periods. This was achieved by calculating the signal-to-noise ratio (SNR) of the wavelet spectrum for all baselines between every combination of remote stations and CS002 (52.915°N, 6.870°E; other core stations were neglected in this step to avoid biasing the results towards the very short baselines within the core) and masking the wavelet spectrum to leave only regions where the median SNR was greater than 3 (corresponding to the 95% confidence level for identifying a signal in the wavelet spectrum; Torrence & Compo, 1998). The spectrum was also masked for all periods below 1 min and for all points within the cone of influence (CoI: the region in which the wavelet coefficients are unreliable due to edge effects from the start or end of the time series). Including the region below a 1-minute period was found in test cases not to provide any reliable fits and so was excluded from the full analysis to save computational time. Within the remaining unmasked regions, for each period the wavelet spectrum was sampled at a time step of one tenth the width of the CoI, which defines the time range across which the wavelet coefficients at a given period are correlated. The width of the CoI Δt for a given period T is given by (Torrence & Compo, 1998)

$$\Delta t = \frac{4\sqrt{2}T}{5}. \quad (1)$$

Samples were taken every tenth of a CoI width to intentionally over-sample at this stage, ensuring that there would be a sample sufficiently close to the time of peak amplitude for any given wave activity.

For some of the sampled times and periods, despite the significant power measured by the wavelet transform, there may not actually be wave activity or it may not be adequately described by the simple plane wave model. In order to remove these fits from further analysis, fits were flagged if they had a reduced chi-squared higher than 5, relative uncertainty in wavelength or amplitude of greater than 50%, estimated measured amplitude on a 30 km baseline below 1 mTECu or estimated

wavelength greater than 1000 km. The minimum amplitude was set based on the known precision of LOFAR dTEC data (Meivius et al., 2016) and the maximum wavelength was set to twice the upper limit estimated in Boyde et al. (2024) based on synthetic dTEC values. The limits on amplitude and wavelength were deliberately set to allow most fits through at this stage, as it was possible to further limit the range of wavelengths and amplitudes considered later in the analysis.

Once the solutions had been calculated and masked, the next step was to identify individual wave “events” from within them. A given wave will show up at multiple times and periods, both due to its own finite duration and range of frequencies and the non-orthogonality of the wavelet basis causing it to be further spread in the wavelet spectrum. In order to determine groups of solutions corresponding to a single wave event, a version of the DBSCAN clustering algorithm (Ester et al., 1996) was applied. This algorithm identifies clusters within data based on the density of points within the N -dimensional space, and also identifies and excludes noise from the identified clusters. It can operate with an arbitrary distance metric, and requires two further parameters to be specified. One of these is the neighbourhood size ϵ and the other is the number of points required within a points neighbourhood to form a cluster $MinPts$.

Points from a single wave event can be expected to be both close in period and time (i.e. close in the wavelet spectrum) as well as having similar fitted wave parameters, specifically azimuth of propagation and phase velocity which should be reasonably stable both with time and across periods. Based on this concept, two distance metrics $r_{spec;i,j}$ and $r_{fit;i,j}$ between a given pair of points i and j were defined as

$$r_{spec;i,k} = \sqrt{\frac{(\log_2(T_i) - \log_2(T_j))^2}{\epsilon_T} + \frac{(t_i - t_j)^2}{\langle \Delta t \rangle_{ij}^2 \epsilon_t}}, \quad (2)$$

and

$$r_{fit;i,j} = \sqrt{\frac{(v_i - v_j)^2}{(\sigma_{v;i}^2 + \sigma_{v;j}^2) \epsilon_v} + \frac{(\theta_i - \theta_j)^2}{(\sigma_{\theta;i}^2 + \sigma_{\theta;j}^2) \epsilon_\theta}}, \quad (3)$$

where t is the time within the observation, $\langle \Delta t \rangle_{ij}$ is the mean CoI width at periods T_i and T_j calculated from (1), v is the estimated phase velocity, θ is the estimated propagation azimuth, σ_v and σ_θ are the respective uncertainties and the ϵ_x parameters are scales for the various distance components. The overall distance $r_{i,j}$ was then taken as the maximum of $r_{spec;i,j}$ and $r_{fit;i,j}$. As discrete wave events within a given observation were generally clearly distinguished visually, the ϵ_x parameters were then determined manually along with the threshold minimum number of points to form a cluster, which will be discussed below. Given the freedom of these ϵ_x parameters within the distance metric, the global neighbourhood size parameter ϵ could simply be set to 1 in this implementation. This was because ϵ defines the threshold value of $r_{i,j}$ required for two points to be considered neighbours, and any given effective size could be obtained holding $\epsilon = 1$ and scaling the ϵ_x parameters.

Given the two part definition of the distance metric in terms of r_{spec} and r_{fit} and the method of sampling the wavelet spectrum space described above, it is possible to define the

maximum possible number of neighbours for a given point in the wavelet spectrum space, which is determined by the scale factors ϵ_T and ϵ_τ . There is no obvious criteria for defining “correct” values for these parameters, so in practice the values were selected based on those which gave reasonable results in test cases, in that regions of obviously consistent wave activity were grouped into clusters without clearly distinct regions being combined together. These suitable values for ϵ_T and ϵ_τ were found to be 0.25 and 0.75, respectively, although in practice the results are not especially sensitive to moderate (i.e. up to a factor of ~ 1.5) changes in either of these parameters. The maximum number of neighbours was found by sampling assuming that the median SNR is greater than 3 for all times and periods and then calculating the number of neighbours for each sampled point assuming that $r_{\text{fit},i,j}$ is less than 1 for all pairs of points. The result of this is shown in Figure 2, and it illustrates why using a fixed value of *MinPts* was found to be unsuitable in this case. At periods above roughly 16 min, the maximum number of neighbours drops substantially due to the presence of the cone of influence. If *MinPts* was set low enough to detect clusters in this region it was found to also detect spurious clusters especially at the lowest periods. To remedy this, the DBSCAN algorithm was adjusted to determine *MinPts* individually for each point considered. The maximum number of neighbours shown in Figure 2 was averaged at each period to obtain a representative value. A point was then considered a core point in DBSCAN if its number of neighbours was 60% or more of the maximum for its period. For the scale factors ϵ_v and ϵ_θ , these were both set to $\sqrt{2}$ to correspond to a “1 sigma” difference in the fit values.

An example of the clustering results is shown in Figure 3. In this observation, three distinct wave events were identified, as highlighted by the coloured borders. As the clustering provides set of points which correspond to the same wave event, each cluster can be used to calculate representative wave parameters for the whole event. First, the fit with the highest amplitude and all other fits which have amplitudes higher than the maximum minus its estimated uncertainty were identified. The mean time and period corresponding to this subset of fits in the cluster defined the points shown as coloured crosses in Figure 3, the representative wave period and time of occurrence for the cluster. The representative wave amplitude and its uncertainty were calculated as the weighted mean and standard deviation of the amplitudes of this subset, with the weights being the inverse of the fit amplitude uncertainty multiplied by the period. The multiplication by the period was necessary due to the reduction in sampling density with increasing period, to avoid biasing the results towards shorter periods. The representative phase velocity was then calculated in the same way, except in this case using the whole cluster not merely the highest amplitude fits. The representative azimuth was calculated by finding the azimuth of a weighted sum of unit vectors along each fitted azimuth direction, again using the whole cluster. This was preferred to simply averaging the fitted azimuths as averaging azimuths runs into issues when the direction is close to the value at which the azimuth values wrap around (i.e. go from 360° back to 0°) and would provide unphysical values in those cases. The 5th and 95th percentiles of period and time within the cluster were also recorded as a measure of how extensive the wave activity was in period and time.




Figure 2. The maximum possible number of neighbours for different locations in the wavelet spectrum. This assumes median SNR > 3 at all points and $r_{\text{fit},i,j} < 1$ for all i and j . The shaded and cross-hatched region is the cone of influence.




Figure 3. An example of the results of the clustering algorithm described above. Shown in the colours are the estimated azimuths as a function of time and period for an observation made on 25th July 2017 (observation ID: L601030). The coloured outlines indicate the three clusters identified within the fit results. The coloured crosses indicate the point within each cluster at which the highest amplitude was estimated. All coloured regions outside the clusters shown were classified as noise. The cross-hatched region shows the cone of influence.

3 Results

The method described above in the previous section identified 7195 wave events within the 2629 h of observations considered, with periods ranging from 1 to 27 min. As shown in Figure 1, the observations covered low to moderate solar activity conditions and primarily low geomagnetic activity. Wave activity was observed for all solar and geomagnetic conditions as well as for all times and seasons.

3.1 Wave occurrence rates

The overall wave occurrence rate was determined by taking the total duration of the wave events in the data and dividing by the total duration of observations, binning by different conditions. The total duration of a given event was defined by the time between the 5th and 95th percentiles of times within the cluster, plus twice the CoI width at the representative wave period. For example, for the wave with a representative period around 10 min in Figure 3 (shown in red), the 5th and 95th percentiles are 13:42:18 and 13:57:40 UT respectively, and the CoI width is 11 min. Combining these gives a duration of around 37.5 min for this wave event. The addition of the CoI width factor was necessary as the period of time in which the wave activity is detected by the wavelet transform will not represent the whole duration of the wave activity. This is particularly clear at longer periods, such as the cluster outlined in blue in Figure 3, where the cluster is bounded on both sides by the CoI and it is unreasonable to assume that wave activity was not present throughout a much greater part of the observation than the cluster covers. Twice the CoI width is arguably a large correction in cases where the cluster is not actually limited by the CoI, but reducing this to adding only one times the CoI width was found to not significantly affect the variations in occurrence rate, only the magnitude of the values, and hence does not change the conclusions that can be drawn about what conditions are more or less favourable for the waves observed by LOFAR.

This definition of the occurrence rate means that an occurrence rate over 100% is possible, which corresponds to multiple waves at different periods being observed simultaneously. The definition also does not include any weighting for the amplitude of the waves. This is because the observed amplitude is dependent not just on the physical properties of the wave in the ionosphere, but also the observing geometry: waves which have their phase fronts aligned with the line of sight will cause larger amplitudes of dTEC perturbations than waves oriented in other directions. In order to avoid biasing the results in this way, all wave events were considered equally for these calculations, regardless of amplitude.

The occurrence rate as a function of UT hour and month is shown in Figure 4, for all wave periods in the top left panel and for subsets of wave periods in each other panel. Focusing on the top left, the overall occurrence rate of the observed wave events peaks between the winter solstice and spring equinox post midday, with a minimum occurrence rate in summer post midday. However, even at this minimum occurrence rate, the rate is still around 100%, indicating that wave activity of some sort is completely ubiquitous in LOFAR observations (note that this still does not mean that waves are guaranteed to be present in every observation as many observations have multiple waves simultaneously observed at different periods).

When the occurrence rate is split by observed period as in the other panels of Figure 4, however, differences in behaviour with period are clearly apparent. For example, the longest periods observed (20–27 min, bottom right panel) show peak occurrence rates around morning and evening in summer, and lower occurrence rates between winter and spring where the overall occurrence rate maximises. The peak between winter and spring only becomes apparent for periods below 15 min (top right and middle panels). This difference in the variation

of occurrence rate for different wave periods indicates that either the sources of waves of different periods are different, or that the propagation conditions that favour AGWs of one period reaching ionospheric heights and being detected by LOFAR are not the same as the conditions favouring waves of another period. It is also clear that the average occurrence rate tends to increase for increasing periods, which is also shown in Table 1. This is despite there being more wave events at short periods, because the events at longer periods generally have longer durations and therefore contribute more to the calculated occurrence rate.

When splitting observations by period, there is a significant problem that must be considered, namely the Doppler shift between the observed period and the real physical period due to the motion of LOFAR's line of sight through the ionosphere. As long as this has a component along the wavevector, this will increase or decrease the apparent velocity and hence period of the wave as seen by LOFAR. This effect can be estimated for a given altitude, and is predominantly <10% for an assumed altitude of 350 km (which will be shown to likely be an overestimate of the true altitude of the waves observed), meaning that for such a coarse binning by period the effect on the conclusions should be minimal. If the altitude of the structures could be estimated, this could be corrected, but this is not practical with the available data from the calibration solutions. The Brunt-Väisälä frequency Ω_B is the natural frequency of vertical oscillations of air parcels in the atmosphere, and typically corresponds to a period of ~ 10 min in the F-region ionosphere (e.g. Snively & Pasko, 2003). Waves with frequencies above Ω_B will be unable to propagate effectively to higher altitudes and so the short period waves observed are likely to be at significantly different altitudes to the longer period waves (which are expected to predominantly exist at F-region altitudes; e.g. Kirkland & Jacobson, 1998; Fedorenko et al., 2011), preventing a single representative altitude being used to correct the Doppler effect. The shorter period waves may be in the E-region, where Ω_B is typically higher than the F-region (e.g. Snively & Pasko, 2003), but waves at these periods have also been observed in the plasmasphere using similar observations from the Very Large Array (Hoogeveen & Jacobson, 1997). Due to these difficulties, the observed period was chosen as a means of distinguishing the observed waves for this analysis.

3.2 Wave parameters

To place the observations made using LOFAR into context with other TID observations, it is also important to consider the range of wavelengths which are observed. Different observations and different analysis methods will be more sensitive to different scales, which makes them complementary to one another but also requires us to consider whether any differences in the observed behaviour are due to real physical differences in the ionosphere-thermosphere system or systematic effects arising from the observing technique. The wavelengths of the waves detected with LOFAR are shown in Figure 5, with the same arrangement of panels as Figure 4. Firstly, this shows that the wavelength cutoff of 1000 km applied in processing the data had a negligible impact, as even for the longest periods observed very few waves are close to the cutoff. The longer periods tend to have larger wavelengths, with the modal wavelength varying between ~ 80 –160 km between different period

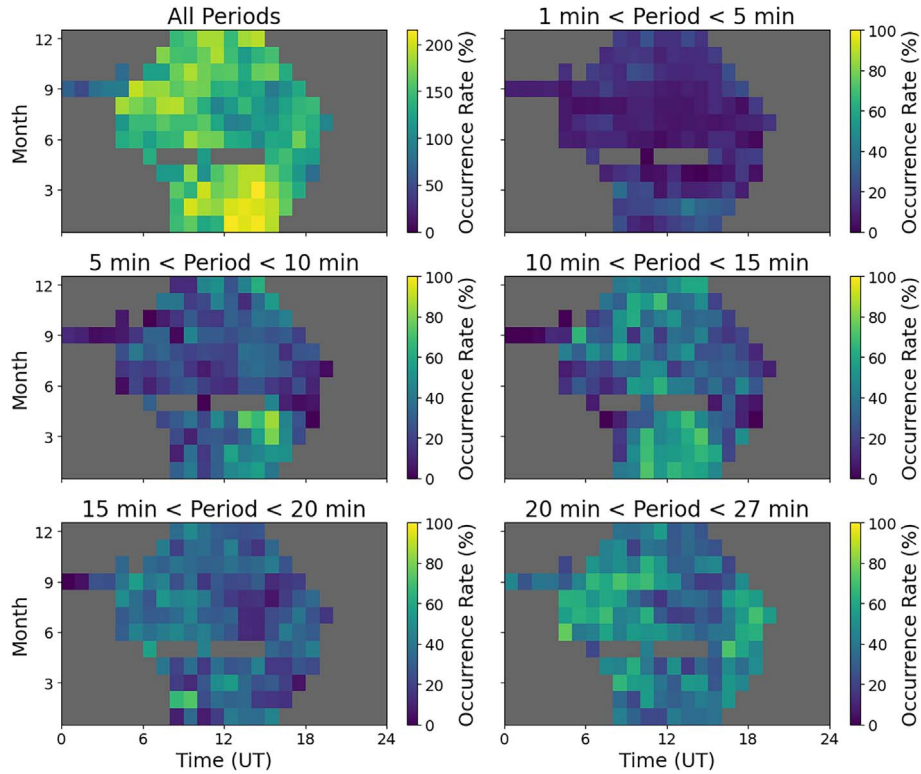


Figure 4. The occurrence rate of wave events as a function of UT hour on the horizontal and month on the vertical. The top left panel includes waves of all periods, other panels show only a subset of observed periods. All panels except the top left use a common colour scale capped at 100%. Only bins with at least 5 h of observations are included, bins with fewer observations are coloured in grey.

Table 1. The number of wave events in each period range used for the analysis.

Period range	No. of waves	Mean occurrence rate
1–5 min	2131	12.9%
5–10 min	1732	27.5%
10–15 min	1386	35.2%
15–20 min	844	30.2%
20–27 min	1102	42.6%
All periods	7195	148.5%

ranges, towards the short end of what is typically considered a medium scale TID (MSTID). The relative lack of waves observed at longer wavelengths is likely due to the reduced sensitivity of LOFAR to wavelengths greater than twice the longest baseline (longest baseline ~ 50 km), meaning that waves of longer wavelength will require higher TEC amplitudes to be detected. Also, as longer wavelengths will more often correspond to longer periods (as indicated by the larger tail of the distribution for periods of 20–27 min in the bottom right panel), many longer wavelength TIDs will have periods too large to be detected with 1-hour observing windows. Shorter wavelengths may also be suppressed to an extent, due to the assumption of an ideal plane wave spanning the whole LOFAR array. For shorter wavelength waves this requires retaining the plane wave assumption over several wavelengths both along and perpendicular to the propagation direction, which is likely to not be the case for some of these waves. However, the peaks

wavelength distributions derived from LOFAR observations are nevertheless similar to those derived from incoherent scatter radar and ionosonde data in Irkutsk by [Medvedev et al. \(2015\)](#), suggesting that there is some physical significance to this wavelength preference.

A subtle difference observed in [Figure 5](#) is the more significant tail of the distribution towards longer wavelengths for the lowest periods (1–5 min, top right) than at other periods. This may be an indication that at least some waves observed at these periods are from a different population to those observed at other periods, possibly being the waves in the plasmasphere reported by [Jacobson & Erickson \(1993\)](#) and [Hoogeveen & Jacobson \(1997\)](#).

The phase velocities of the detected waves are shown in [Figure 6](#) in the same format as the wavelengths in [Figure 5](#). They show that the majority of the waves have velocities between ~ 50 – 250 ms^{-1} , as is typical for small- to medium-scale TIDs (e.g. [Hunsucker, 1982](#); [Hocke & Schlegel, 1996](#)). The exception to this are the short period waves with periods below 5 min which show a much broader velocity distribution extending well above 1000 ms^{-1} . This is further evidence to support the idea that these waves are a distinct population to the longer period waves ($> \sim 5$ – 10 min), and consistent with the short period waves being located in the plasmasphere where the higher effective scan velocity of the line of sight can cause such large apparent velocities ([Jacobson & Erickson, 1993](#); [Hoogeveen & Jacobson, 1997](#)). The differences between the short period waves and the others are explored in more detail in [Section 4.3](#).

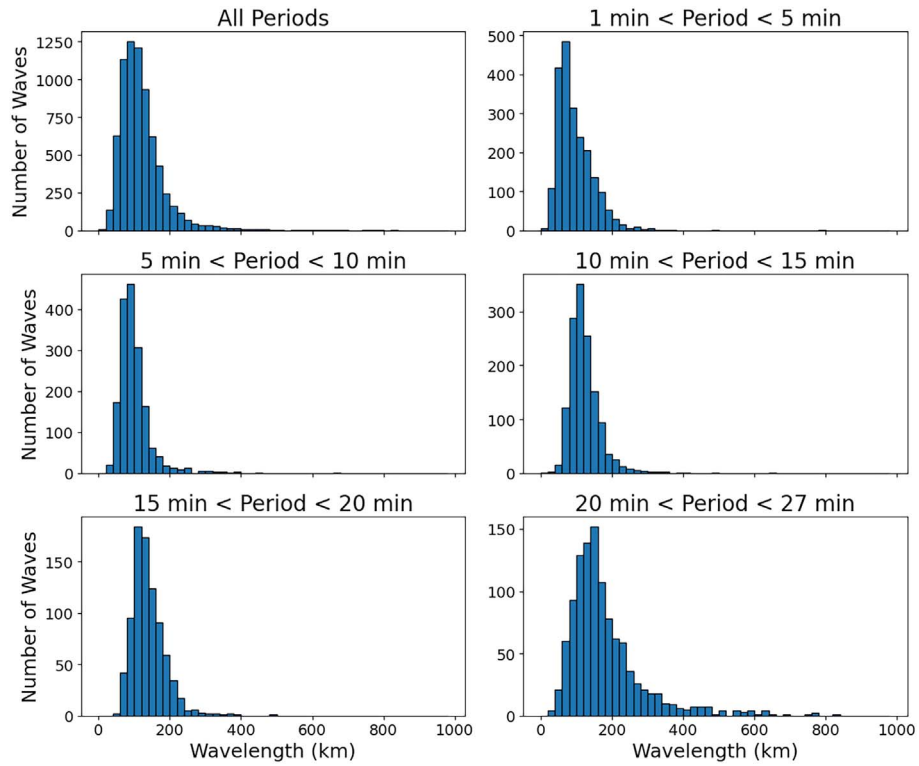


Figure 5. The distribution of observed wavelengths, with the panels corresponding to those in Figure 4. Note each panel uses a different scale for the vertical axis, but all are on the same horizontal axis scale.

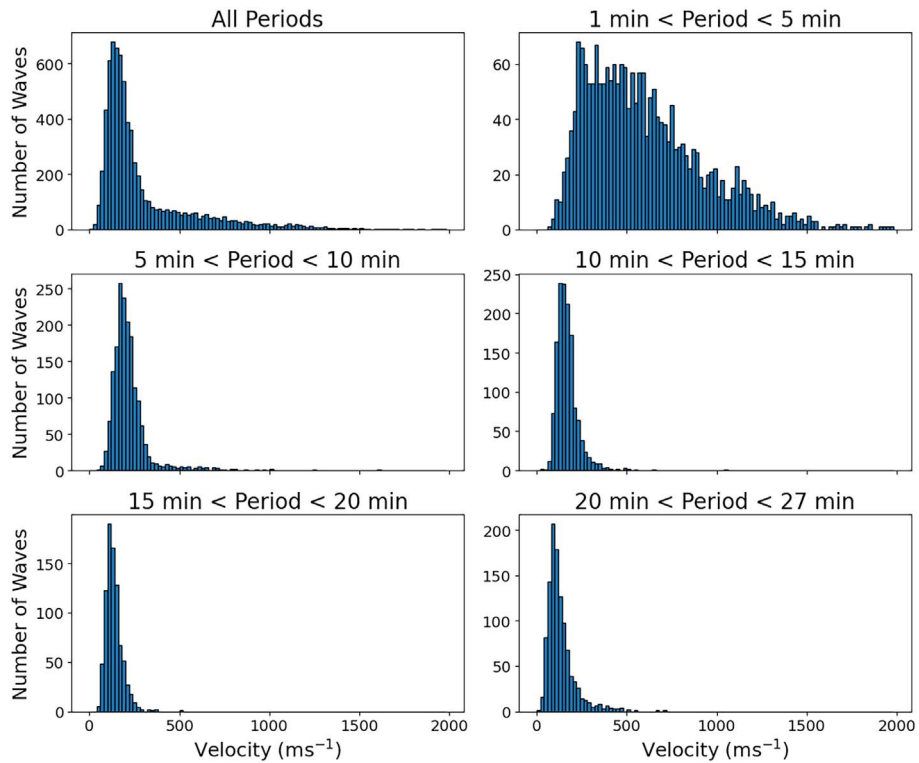


Figure 6. The distribution of observed phase velocities, with the panels corresponding to those in Figure 4. Note each panel uses a different scale for the vertical axis, but all are on the same horizontal axis scale.

3.3 Propagation directions

As mentioned in Section 1, the direction of AGWs, and hence the TIDs cause, is strongly influenced by the neutral wind direction, with waves propagating against the background wind able to propagate to higher altitudes and others filtered out (Cowling et al., 1971). No direct observations of the neutral winds in the thermosphere during our observations were available, but the Horizontal Wind Model 2014 (HWM14: Drob et al., 2015) can provide a climatological estimate of the winds. This model has substantial uncertainties in predicting the wind at a single time (root mean square error of $\sim 50 \text{ ms}^{-1}$ in each component, see Table 2 of Drob et al., 2015), but provides a good description of the systematic seasonal and diurnal variation in the wind. The altitude range at which the filtering occurs is uncertain, so the observations were compared to the modelled wind directions at various altitudes to determine if there is a relationship. The winds were calculated for the latitude and longitude of the ionospheric pierce point (IPP, the intersection between the line of sight and a specified altitude shell) at 350 km from station CS002 at the time of the individual wave event. Choosing different altitudes for the assumed IPP was found to have negligible impact on the results (differences of a fraction of a degree in latitude and longitude) as observations are taken at high elevation. Overall the IPPs of wave events ranged between approximately $0\text{--}14^\circ\text{E}$ and $52.5\text{--}60^\circ\text{N}$, with the majority between $3\text{--}9^\circ\text{E}$ and $52.5\text{--}54^\circ\text{N}$. HWM14 also has dependency on geomagnetic activity through the ap index, although the impacts of this on wind direction in the mid-latitudes are typically very small, especially considering the generally low geomagnetic activity during the LOFAR observations shown in Figure 1. Figure 7 shows the observed wave directions for periods between 20 and 27 min in the top left panel, along with the corresponding HWM14 anti-wind directions at a range of altitudes from 120 to 280 km in the other panels, as a function of time and day of year. Simply by inspection of these plots, it is clear that the agreement between observed wave direction and modelled anti-wind direction is reasonable for altitudes above ~ 200 km and that below that it begins to break down. This strongly indicates that the TIDs observed by LOFAR are in the F-region of the ionosphere.

While the simple visual inspection of Figure 7 shows that the observed wave directions do correspond well to expected anti-wind directions in the thermosphere, and give an approximate sense of the relevant altitude, a more quantitative picture is provided by Figure 8. The left panel shows the weighted standard deviation (weighted by the modelled wind speed) of the difference between observed wave direction and modelled anti-wind direction for altitudes between 80 and 300 km with an altitude step of 10 km. The standard deviation was weighted by the modelled wind speed, to reflect both the fact that the wind filtering effect is stronger for higher wind speeds and that when the expected wind speed is low the uncertainty in the wind direction is likely to be higher. This illustrates what was apparent from Figure 7, with the lowest values of the standard deviation being at altitudes of 200 km and above. The standard deviation minimises at 220 km, and the distribution of offsets in time and day of year relative to the winds at this altitude is shown in the right panel of Figure 8. This shows that although the winds at this altitude provide a reasonable explanation for the observed variation of wave propagation directions, there

are still unresolved aspects such as the overwhelmingly negative offsets for day of year 150–250 (roughly June–August) and UT < 10 h. The systematic variation of the offset with time and day of year, rather than purely random noise, suggests that there is a systematic effect not accounted for in the analysis here such as variation in the optimum altitude with time and season. Such a variation could be explained by the AGW behaviour being primarily determined by pressure level rather than altitude, for example, but this is not explored further here.

It is important to emphasise here that the filtering of AGWs by the wind does not happen at a single altitude, but is a process that is active throughout their vertical propagation from the lower atmosphere up to the thermosphere. No attempt is made here to estimate the integrated effect of this process on the observed waves, as the off-vertical propagation of AGWs and their finite vertical velocities cannot be accurately accounted for with the data available in this study. The optimum altitude estimated here is also not a reliable measure of the altitude at which the TID amplitude maximises or of the “effective altitude” of the integrated perturbations detected by LOFAR, but simply reflects the altitude at which the modelled winds best describe the observed seasonal and diurnal variations in wave directions.

Based on Figure 8, an altitude of 220 km can be taken as reasonably representative of the altitude relevant to the longest period waves in the data. The anti-wind direction at this altitude was then compared to the observed wave direction across the other period bins to investigate any systematic differences in propagation direction with period, which is shown in Figure 9. We note at this point that given the typical Brunt-Väisälä period at 220 km altitude is ~ 10 min, we would not expect waves with periods below 10 min to show strong agreement with the winds at this altitude. However, the comparison is carried out at all periods for completeness as it nevertheless highlights differences between the propagation directions of waves of different periods with respect to one another, and comparison of the short period waves to the winds is still physically reasonable when considering other altitudes, such as those of the E-region where the Brunt-Väisälä period is often substantially lower (e.g. Snively & Pasko, 2003). The vertical axis on each panel of Figure 9 shows the offset between the wave direction and anti-wind direction, with negative values corresponding to anti-clockwise rotation of the wave with respect to the anti-wind direction as in Figure 8, while the horizontal axis distinguishes the period bins. The top left panel of Figure 9 shows the results including all wave observations, while the other three panels show subsets of time and season. The panels labelled Summer (top right and bottom left) include data with day of year between 90 and 270 inclusive, with pre reversal including all waves before 15:30 UT (2,541 waves) and post reversal including all after 15:30 UT (857 waves). This separation was made due to the clear reversal of both observed wave direction and modelled anti-wind direction around this time shown in Figure 7, and Figure 9 shows a clear change in behaviour corresponding to this time. The panel labelled Winter (bottom right) includes all observations with day of year less than 90 or greater than 270 (3,797 waves).

Figure 9 shows clear period dependence in the wave propagation directions and how they relate to the neutral winds. In all panels, there is a clear tendency for the waves to rotate with respect to the 220 km winds as a function of period (except for

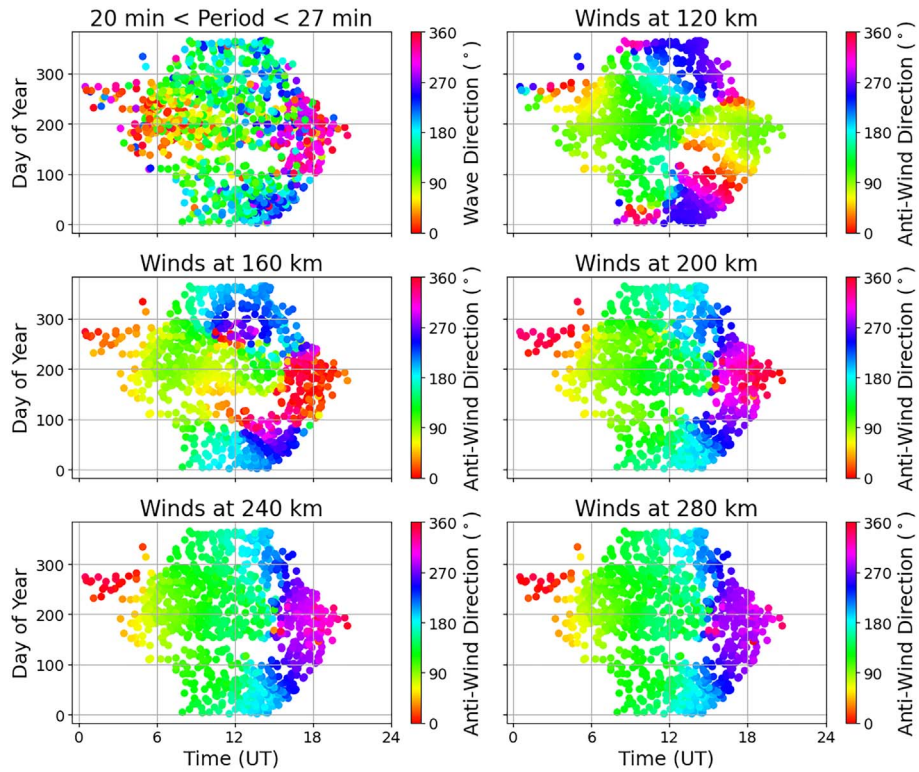


Figure 7. The variation of the observed wave direction for periods 20–27 min with time and day of year (top left), and the modelled anti-wind direction from HWM14 at a range of altitudes (other panels).

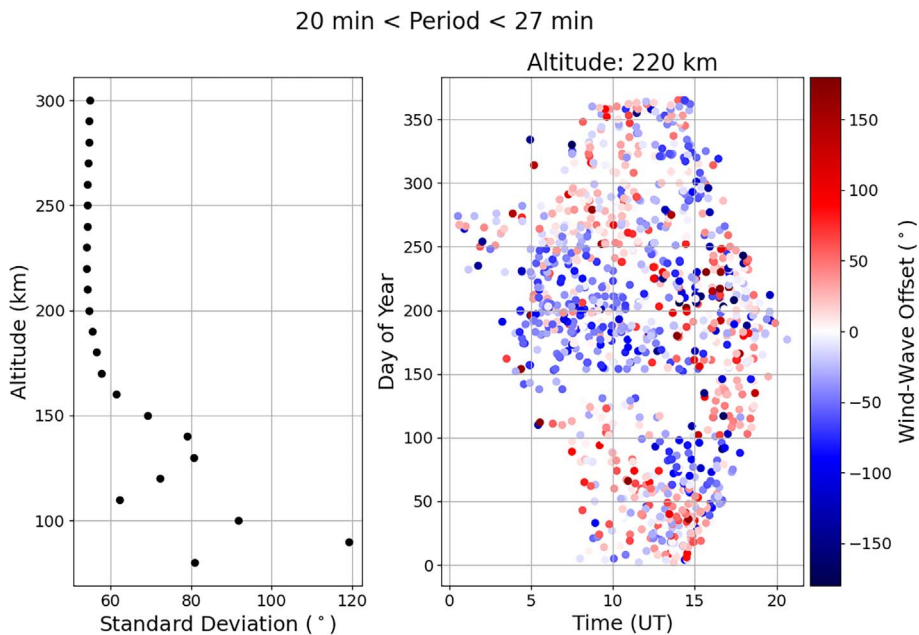


Figure 8. The left panel shows the weighted standard deviation of angular offsets between HWM14 predicted anti-wind direction and observed wave direction for periods between 20–27 min as a function of altitude, while the right panel shows the angular offsets as a function of time and day of year for an altitude of 220 km, where the standard deviation is minimised. Negative offsets correspond to the wave being rotated anti-clockwise with respect to the anti-wind direction, and positive offsets correspond to clockwise rotation. The calculation of the standard deviation is weighted by the wind speed.

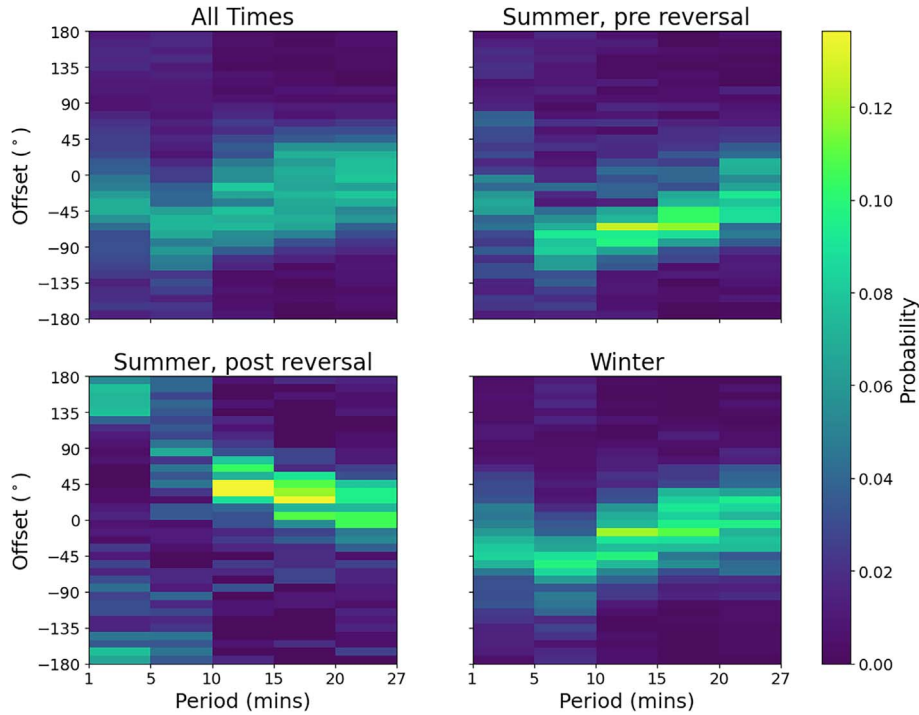


Figure 9. The distribution of wave-antiwind offsets for the different period ranges, each panel showing a different range of time and day of year. Summer is defined as day of year between 90 and 270 inclusive, whereas winter includes all days after day of year 270 and before day of year 90. Pre and post reversal refer to before and after 15:30 UT respectively. Winds are taken at an altitude of 220 km. The colours indicate the proportion of waves within the given period bin which fall into each 10° wide azimuth bin.

the 1–5 min bin which breaks this trend), but the sense and magnitude of this rotation is not constant. The clearest example of this is that for both “Summer, pre reversal” and “Winter”, the rotation of the shorter periods with respect to the wind is anticlockwise (negative offset), whereas for “Summer, post reversal” it is clockwise. Comparing the “Summer, pre reversal” and “Winter” panels more closely, it is apparent that despite the rotation having the same sense, the magnitude is different, with the 10–15 min period bin peaking around $10\text{--}50^\circ$ for “Winter” and around $50\text{--}90^\circ$ for “Summer, pre reversal”.

Given these differences in behaviour for different periods shown in Figure 9, the analysis shown in Figure 8 was repeated for the other period ranges to determine if different altitudes provide better agreement for the different bins. The results of this analysis are summarised in Figure 10, showing that the general trend is for the shorter periods to show better agreement with the winds at slightly lower altitudes (again periods of 1–5 min provide an exception to this). Starting from the longest periods in the data, the optimum altitude, at which there is the best match between modelled winds and observed waves, decreases with decreasing period as shown in the right panel. This starts to break down for the 5–10 min period range, where the magnitude of the standard deviation also increases significantly as shown in the left panel. For the shortest periods of 1–5, the waves do not show any sign of tracking the diurnal rotation of wind direction and so there is no clear optimum altitude (see Fig. 14).

The evolution of the observed propagation directions with time and day of year, along with the corresponding anti-wind direction at the optimum altitude, are shown in Figure 11 for

the period ranges above 10 min (the shorter periods are shown and discussed in Sect. 4.3). The general agreement with the solar heating driven wind variations for all periods, especially in terms of the diurnal variation, is clear. The differences between the winds for the three different altitudes shown in Figure 11 are relatively small, which explains the relatively small variations with altitude of the standard deviations shown in Figure 10. The greatest differences between modelled winds at different altitudes, such as the roughly northward propagation around 12 UT and day of year 275 in the bottom-right panel of Figure 11, correspond to small wind velocities ($< \sim 20 \text{ ms}^{-1}$) and hence have limited impact on the propagation of waves with typical phase velocities of $\sim 50\text{--}250 \text{ ms}^{-1}$. In terms of differences between the observed wave directions and the modelled winds, there is a clear tendency in the summer mornings for the waves to be rotated anticlockwise with respect to the winds, becoming clearer at the shorter periods. Other authors have suggested a time lag between the wind direction and observed wave direction from both observations (e.g. Waldock & Jones, 1986; Xu et al., 2024) and theoretical calculations (e.g. Cowling et al., 1971) due to the time taken for AGWs to propagate vertically from the altitude at which wind filtering is most effective to the altitude at which they are observed as TIDs. However, the near immediate response to the change in direction in the afternoon (around 15–16 UT) in summer shows that the observed waves are able to track changes in the wind very rapidly.

As mentioned previously, the wind filtering explanation for the observed wave behaviour suggests that the filtering effect should be stronger for higher wind speeds relative to the wave

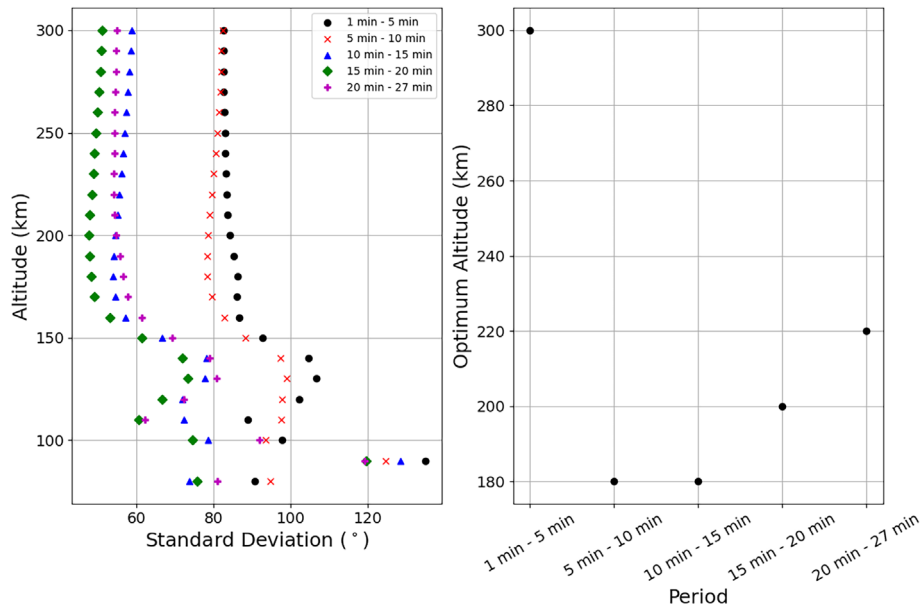


Figure 10. The weighted standard deviation of angular offsets between HWM14 predicted anti-wind direction and observed wave direction as a function of altitude for the various period ranges (left panel) and the altitude at which the standard deviation minimises for each period range (“optimum” altitude, right panel). Standard deviation is weighted by wind speed as in Figure 8.

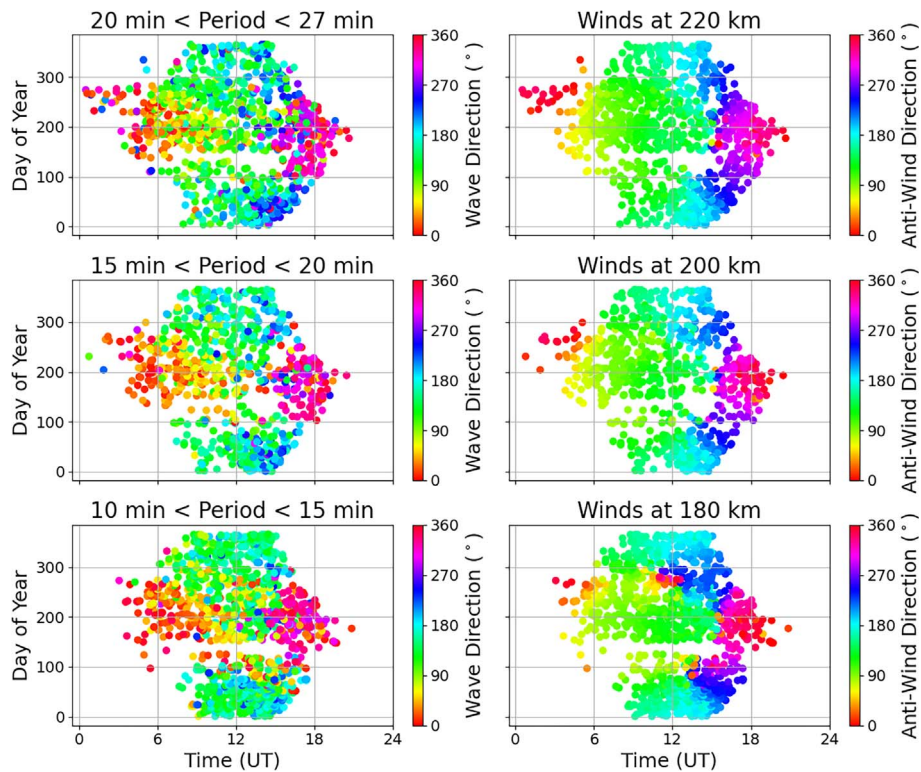


Figure 11. The left column shows observed wave propagation directions and the right column shows HWM14 modelled anti-wind directions at the corresponding times for the optimum altitude calculated in Figure 10. Each row corresponds to a different period range as indicated on each of the left hand panels.

phase speed. In order to test whether this is supported by the data obtained using LOFAR, Figure 12 shows the distribution of wave-antiwind offsets as a function of the ratio of modelled

wind speed to observed wave speed (“velocity ratio”) in the left panel and the standard deviation of the offsets in each velocity ratio bin in the right panel. In order to carry out this part of the

10 min < Periods < 27 min, Altitude 200 km




Figure 12. The left panel shows the distribution of wind-wave offsets as a function of the ratio of HWM14 modelled wind speed divided by observed wave phase speed (“velocity ratio”), with the wind taken at an altitude of 200 km and wave periods from 10 to 27 min considered. The right panel shows the standard deviation of offsets in each velocity ratio bin of the left panel.

analysis, waves of all periods from 10 to 27 min were considered together to ensure there were enough waves in each velocity ratio bin. This is particularly necessary given the fact noted above that HWM14 is only reliable as a climatology and will have substantial uncertainties for any individual wave event. The winds were considered at an altitude of 200 km as this was the “optimum” altitude for all these waves considered together derived in the same way as shown in Figure 10 (not shown). Figure 12 shows that for lower values of the velocity ratio (i.e. when wind speed is significantly lower than the wave phase speed) the agreement between wave and anti-wind direction is significantly weaker than for higher velocity ratios. For velocity ratios greater than ~ 0.5 there is no discernible variation in how well the modelled wind directions predict the wave directions, which may reflect a physical limit but could also simply reflect the limitation of the accuracy of HWM14.

3.4 Wave amplitude

Although the individual wave amplitudes should be treated with caution due to the impact of observing geometry, as discussed previously, the average wave amplitude under different conditions can provide useful information. Figure 13 shows the mean wave amplitude as a function of time and solar activity for all waves (top left) and subdivided into the different period bands in the other panels. As expected, it shows that higher amplitudes were observed under higher solar activity and post-midday when background ionospheric densities are highest. Comparing the different period ranges, there is little difference in the observed amplitudes for the highest bands (10–15 min, 15–20 min and 20–27 min) whereas the shorter periods show significantly lower amplitudes, especially in the 1–5 min period band. This supports the interpretation that the short period

(< ~ 10 min) waves are dominated by a distinct population to the longer period waves as suggested by the difference in propagation directions relative to the neutral wind.

4 Discussion

4.1 Occurrence rates and possible wave sources

The occurrence rate of daytime MSTIDs in the mid-latitudes has been studied by a range of authors, with some generally consistent results. The main reported pattern is a maximum of occurrence in winter and a minimum in summer (e.g. Waldock & Jones, 1986; Ding et al., 2011; Otsuka et al., 2013; Sivakandan et al., 2021), which has also been reported for AGWs from in-situ satellite measurements (e.g. Park et al., 2014). This result is not found universally however, with Negrea et al. (2018) reporting roughly equal peaks in activity in both summer and winter. Breaking this down to sub-seasonal timescales, Friswell et al. (2016) found that MSTID activity over North America tracked variations in the strength of the polar vortex, with the winter peak corresponding to the period in which the polar vortex is fully developed and decreasing activity after the solstice as the vortex breaks down. The diurnal variation is somewhat more complicated, with some authors reporting a single peak around midday or soon after (e.g. Ding et al., 2011; Otsuka et al., 2013) and others reporting peaks around sunrise and afternoon/sunset (e.g. Zalozovski et al., 2021). Some of this difference may arise from seasonal differences in the diurnal pattern, as suggested by Jacobson et al. (1995), who reported a single post midday peak in winter and at the equinoxes as opposed to two distinct peaks in morning and evening in summer.

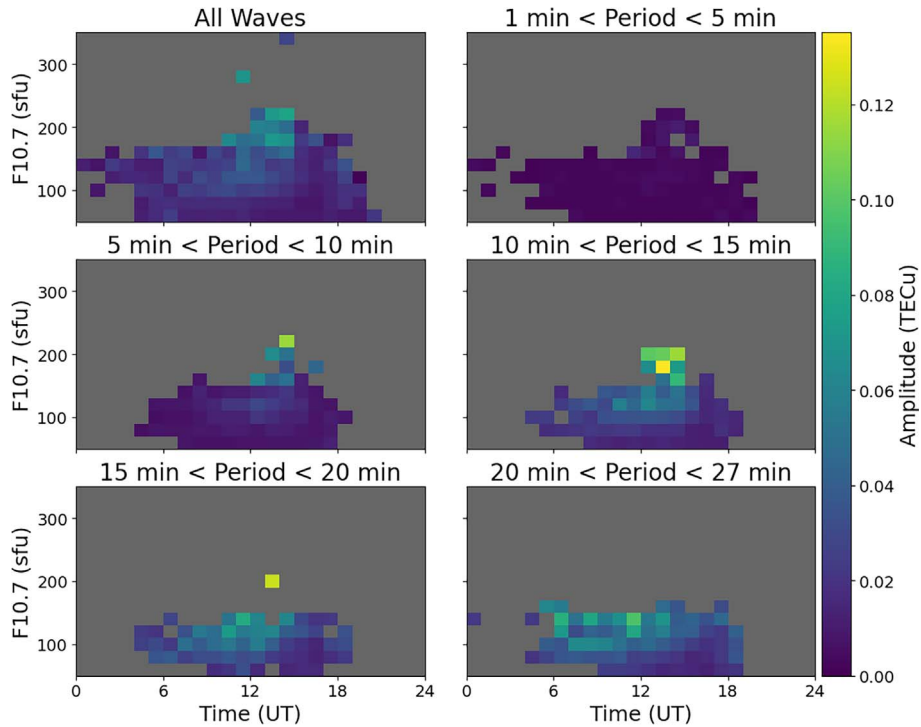


Figure 13. The mean wave amplitude as a function of UT on the horizontal axis and F10.7 solar flux index on the vertical axis. The panels correspond to the same periods as the panels in Figure 4. All panels use the same colour scale. Bins are 1 h by 20 sfu. Only bins with at least five wave observations are included, all others are coloured in grey.

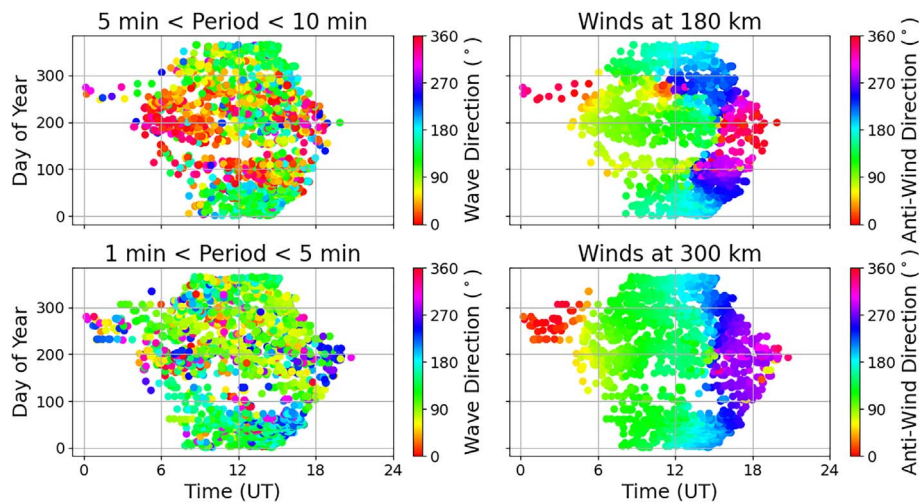


Figure 14. The left column shows observed wave propagation directions and the right column shows HWM14 modelled anti-wind directions for the optimum altitude calculated from Figure 10 at the corresponding times. Each row corresponds to a different period range.

Comparing these to the results from LOFAR shown in Figure 4, there are broad similarities but some subtle differences. The overall occurrence rate (top left, with the caveat that this likely includes some non-ionospheric waves) does show a peak in winter and a minimum in summer as is generally reported. The diurnal behaviour is also broadly consistent with the pattern established by Jacobson et al. (1995) with two distinct peaks in summer and a single peak in other seasons

(although the second peak identified by Jacobson et al. (1995) maximises around 20-23 LT, it becomes apparent around 16-17 LT similar to the observations made by LOFAR). Once this is broken down by period in the other panels of Figure 4, a slightly more complicated picture emerges. The winter peak is prominent for periods of 5–15 min, whereas for periods of 15–20 min there is no discernible seasonal variation. For periods of 20–27 min there are two peaks in summer, one in the

morning and another in the evening close to sunset, although it is possible that this double peak is also present in winter but cut off by the limited time coverage of the observations in winter.

These differences with period may suggest an explanation for some of the inconsistencies between the TID occurrence rates reported by other authors. Different observing techniques will preferentially detect waves with different parameters (e.g. period, wavelength) which can lead to different conclusions being drawn. For example, [Otsuka et al. \(2013\)](#) used a relatively high amplitude threshold of 0.2 TECu to identify TIDs in GNSS TEC data for their analysis (contrast with the ~ 0.001 TECu threshold used here) which suggests that their conclusion of near zero TID occurrence in Summer may instead reflect lower amplitudes. Due to the broader geographical coverage of GNSS receivers and continuous observations, however, [Otsuka et al. \(2013\)](#) were able to extend their analysis to much larger wavelengths (up to 1500 km) and longer periods (10–60 min) than is possible with the LOFAR dTEC solutions. This emphasises the value of complementary instrumentation and analysis techniques to provide a complete picture of TIDs, as well as the need to account for the different biases when analysing and comparing the results.

The suggestion that reported differences in occurrence rates from different authors are at least somewhat due to variations in amplitude is supported by the results of [Ding et al. \(2011\)](#) and [Jacobson et al. \(1995\)](#), who report that TID amplitudes maximised in winter at the same time as the reported maximum occurrence rate. The variation in amplitude also has a more direct impact on the comparisons as some authors discuss an occurrence rate similar to this work (e.g. [Otsuka et al., 2013](#)), whereas others discuss an “activity level” which in effect combines the occurrence rate variations and amplitude variations into a single measure (e.g. [Frissell et al., 2016](#); [Negrea et al., 2018](#)).

Overall, the generally high occurrence rates reported here suggest that the dominant drivers of the observed waves are not sporadic impulsive events such as geomagnetic disturbances, earthquakes and major tropospheric storms. Rather, it points to drivers such as jet streams and the polar vortex ([Buss et al., 2004](#); [Sato & Yoshiki, 2008](#)) as more likely candidates for the consistent generation of AGWs which then propagate to ionospheric altitudes.

4.2 Period dependence of propagation direction and inferred altitudes

Some authors have previously reported variations in TID propagation with height or period. For example, [Negrea et al. \(2016\)](#) found that the propagation direction of observed TIDs in ionosonde data from October 2013 in the North American mid-latitudes rotated clockwise with increasing altitude during the daytime. However, this was only clearly observed for the longest periods in their data (frequencies < 1 mHz, corresponding to periods $> \sim 17$ min) with shorter periods showing much weaker directional preference, and is limited by averaging all daytime observations together thereby removing any information on variations with local time. This absence of local time information prevents comparison with neutral wind directions as these vary substantially with local time. Subsequently, using roughly 3 years of data from the same location [Negrea et al. \(2018\)](#) showed that the activity of shorter period waves was higher at slightly lower altitudes, which was apparent for

frequencies up to ~ 1.5 mHz corresponding to periods above ~ 11 min. This is consistent with the results presented here showing better agreement between wave directions and modelled neutral winds at lower altitudes for lower periods, and the apparent break down of this relationship for periods below ~ 10 min.

On a case study basis, a further example is provided by [Crowley & Rodrigues \(2012\)](#) who showed TIDs with ~ 15 – 30 min periods displaying anticlockwise rotation with increasing period in data from HF Doppler sounding (see their Figure 7, red box labelled “Medium Scale”). These observations were made between 22–01 LT in the North American sector, meaning they fall outside the geographical and time of day coverage of the present study, but they provide further evidence that this rotation of TIDs with period is detected by a variety of techniques. In this case, it is also interesting to note that as [Crowley & Rodrigues \(2012\)](#) used HF Doppler sounding, all measurements were obtained at the same altitude. This suggests that the observed rotation with period in the data obtained using LOFAR may not be entirely due to the waves peaking in amplitude at different altitudes, but rather the neutral winds at different altitudes having stronger effects on waves of different periods, with these effects persisting as they propagate to higher altitudes.

Extensive theoretical calculations on the vertical propagation of AGWs have been carried out by [Vadas \(2007\)](#), showing the dissipation altitudes (altitude at which the AGW momentum flux maximises) for various AGW parameters. For the range of wave parameters observed here (periods up to 27 min, velocities of 50 – 250 ms^{-1} , wavelengths of 50 – 250 km) the dissipation altitude would be expected to increase with decreasing period (see their Figs. 6 and 7, increasing from ~ 200 km for 20–27 min waves to ~ 225 – 250 km for 10–15 min for moderate solar activity). This seems to run counter to the observations from LOFAR presented here which show wave propagation directions at shorter periods better explained by wind directions at lower altitudes. One possible factor that could explain this discrepancy is that the vertical wavelength of the shorter period waves reported here would be much greater than the longer period waves ([Vadas, 2007](#)), potentially making the shorter period waves more sensitive to the winds at lower altitudes. With this explanation, the observations at all periods could be dominated by contributions from the same altitude (e.g. the altitude of the F2 ionospheric density peak) rather than arising from different altitudes determined by individual wave dissipation altitudes. However, this explanation would not explain the direct observations of shorter period waves being more significant in the ionosphere at lower altitudes made by [Negrea et al. \(2018\)](#).

Regarding the specific altitudes at which the waves observed by LOFAR best correspond to the modelled neutral winds, previous authors have suggested similar values for the altitude of waves observed in TEC. For example, [Kirkland & Jacobson \(1998\)](#) used a radio interferometer in New Mexico which observed signals from two closely spaced geostationary satellites to deduce the effective altitude of the observed TEC perturbations by parallax, and found typical altitudes around 180 ± 40 km, comparable to the range of 180 – 220 km found by matching the wind directions on a statistical basis here. They also reported that events with higher amplitude were biased towards higher altitudes, although still mostly within the

standard deviation of the overall altitude distribution. The data reported by Kirkland & Jacobson (1998) covered only October to March, and they do not provide information on the typical wave parameters (e.g. period, wavelength) in their observations which limits the ability to make more detailed comparisons.

The relationship between the different period waves observed by LOFAR and the neutral winds at different altitudes, and the uncertainties in the underlying physical mechanism, suggest a useful direction for future work. Modelling work to investigate the AGW propagation under the conditions relevant to the LOFAR observations, case studies incorporating data from other instruments such as ionosondes which can provide altitude resolved measurements, or a combination of these approaches, may enable the causes of this relationship to be determined.

4.3 Short period waves: possible plasmaspheric structures

One notable aspect of data obtained from radio interferometers observing natural radio sources compared to other techniques for observing TIDs is their capability to detect plasmaspheric as well as ionospheric disturbances. GNSS satellite orbits are typically too low to observe most of the plasmasphere, and while satellites in higher orbits such as geostationary satellites will observe a plasmaspheric contribution it is difficult to disentangle from ionospheric signals as the ionospheric variations will typically be of much higher magnitude. However, due to the motion of the line of sight to an astronomical source relative to the plasmasphere having a relatively high effective velocity compared to that in the ionosphere, the Doppler shift effect can shift the apparent period of plasmaspheric disturbances to short periods where ionospheric disturbances are minimal (Jacobson & Erickson, 1993). Previous observations using the Very Large Array (VLA: Thompson et al., 1980) have found plasmaspheric disturbances in the short period range ($< \sim 5$ min) with a strong preference for propagation towards magnetic East (Jacobson & Erickson, 1992, 1993). Later investigation showed that these structures are field-aligned plasmaspheric structures which actually tend to move westwards, with the apparent eastward propagation being induced by the motion of the line of sight (Hoogeveen & Jacobson, 1997).

To investigate the possibility of field-aligned plasmaspheric structures being included in the short period waves shown in Figure 14, for each wave event the source azimuth and elevation were used to calculate the projection of the magnetic field into the ground plane for various altitudes along the line of sight to determine if there is an altitude for which the geomagnetic field is perpendicular to the observed propagation (as was done by Hoogeveen & Jacobson (1997), see Appendix for details of these calculations). Altitudes up to $10R_E$ were considered for this, using the International Geomagnetic Reference Field (IGRF: Alken et al., 2021) to provide the geomagnetic field orientation. As this is a model of the core field, it does not include contributions from current systems in the ionosphere or magnetosphere which will become more significant at greater altitudes meaning its accuracy at the highest altitudes considered here is doubtful. However, it will be shown that the altitudes of alignment are predominantly much lower than $10R_E$ ($< \sim 2R_E$, see Fig. 15). At these altitudes the IGRF is a more reliable model of the field than at $10R_E$, and hence restricting the altitude

range further than was done here would in fact strengthen the conclusions drawn rather than weaken them.

Naturally, even if no such field-aligned waves exist in the data some waves would happen to align with the magnetic field regardless, purely by chance. To test if this is the case, a Monte Carlo analysis was also carried out where 20 random directions were generated corresponding to each wave event and each compared to see if they were field-aligned. This Monte Carlo analysis used three different methods for generating random directions: uniform (i.e. completely random) directions, normal distribution around the HWM14 anti-wind direction at 180 km with 30° standard deviation (to represent F-region propagation), and normal distribution around the HWM14 anti-wind direction at 100 km with 30° standard deviation (to represent E-region propagation).

In terms of the number of field-aligned events, for periods of 1–5 min there were 1,028 out of a total of 2,131. From the Monte Carlo results the expected values were estimated as 855, 760 and 756 for the uniform, 180 km and 100 km distributions respectively. This indicates a small but consistent increase in the number of field-aligned waves relative to what might be expected to occur by chance. As well as the number of events, the altitudes at which the field-alignment was found for the real data compared to the Monte Carlo results (shown in the top-left panel of Figure 15) supports the idea that these are field-aligned waves in the plasmasphere. Relative to the Monte Carlo results, the altitudes of field-alignment for the real data are skewed lower, generally at altitudes below $1R_E$, which is expected given that at lower altitudes plasmaspheric densities will be higher and hence perturbation amplitudes are more likely to be high enough to be detected. This emphasises the point that if the altitude range considered were further restricted due to the limitations of using the IGRF at large distances from the Earth, more of the Monte Carlo results would be excluded than the observed waves, strengthening the inference that the observations contain real field-aligned plasmaspheric structure.

Comparing the wavelengths of the field-aligned and non-field-aligned waves in the data (bottom-left panel of Fig. 15) also suggests that they come from different populations, with field-aligned waves having markedly shorter wavelengths. This wavelength difference indicates that the plasmaspheric field-aligned disturbances cannot explain the stronger high wavelength tail for the waves at 1–5 min periods that was noted in Figure 5. Finally, the occurrence of field-aligned events in time and season (top-right panel of Fig. 15) shows a relatively uniform occurrence of field-aligned events across observing conditions, whereas the non-field aligned events cluster at certain times (middle-right panel of Fig. 15). For the Monte Carlo results, using HWM14 winds to constrain expected wave directions causes very clear clustering in time and season for the field-aligned events (when winds are predominantly zonal, as illustrated in the bottom-right panel of Fig. 15) which is not observed in the real data. The same analysis as shown in Figure 15 was also carried out for periods of 5–10 min (not shown) but did not show any evidence of an equivalent population in that period range, with fewer events than predicted by the Monte Carlo simulations and no preference for lower altitudes or significant bias in wave parameters between field-aligned and non-field aligned events.

These results suggest that at least for periods below 5 min there is a population of field-aligned plasmaspheric structures contributing to the waves observed with LOFAR, broadly

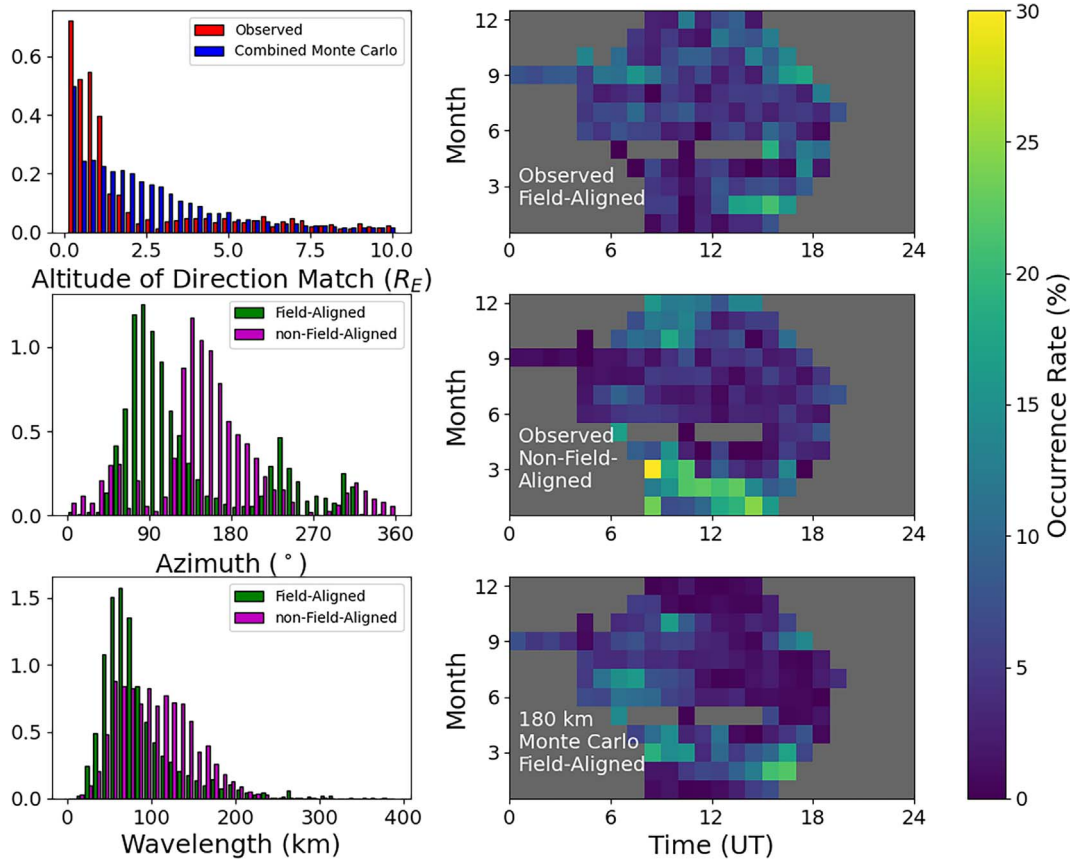


Figure 15. The results of examining the field-alignment of the observed waves with periods between 1–5 min. The top-left panel shows the altitude (in units of Earth radii R_E) at which the propagation direction matched the field orientation for the observed waves (red) and the combination of all three Monte Carlo tests (the Monte Carlo results are combined for clarity as they were all consistent with one another). The middle-left panel shows the observed azimuth of propagation for field-aligned (green) and non-field-aligned (magenta) waves, while the bottom-left panel compares observed wavelengths for the same groups. Each histogram is independently normalised, and the vertical axis gives the relative probability in arbitrary units. The right hand panels show occurrence rates as a function of time and month for the observed field-aligned waves (top-right), observed non-field-aligned waves (middle-right) and Monte Carlo results using 180 km HWM14 winds (bottom-right).

consistent with those observed by [Hoogeveen & Jacobson \(1997\)](#). The ratio of field-aligned events to non-field aligned events is higher than reported from the VLA, with [Hoogeveen & Jacobson \(1997\)](#) reporting 4,581 of 14,017 (32.7%) compared to 1,028 of 2,131 (48.2%) in the data from LOFAR, possibly due to the different geographical locations and different sources observed. The wavelengths and velocities of the field-aligned disturbances are also broadly similar between the two datasets. However, the non-field-aligned waves show clear differences, with strong preference for certain azimuths in the LOFAR data (see middle-left panel of [Fig. 15](#)) compared to relatively uniform azimuth distribution reported in [Hoogeveen & Jacobson \(1997\)](#) (see their [Fig. 3](#)). Their field-aligned waves also constitute a very distinct peak in the azimuth distribution, whereas the field-aligned waves seen here are much less prominent in the overall azimuth distribution. This may be due to different field geometry probed by LOFAR compared to the VLA due to their different locations and different sources observed, which may spread out the range of azimuths at which

field-alignment is possible for the LOFAR observations compared to those of [Hoogeveen & Jacobson \(1997\)](#) and hence suppress the sharp peak around magnetic East.

While this indicates an explanation for some of the observed short period waves, the majority of them are still unexplained. Even with the field-aligned waves removed, there is no apparent match with the climatological neutral winds at any altitude, although the magnitude of the standard deviation is somewhat reduced compared to [Figure 10](#) (a minimum of $\sim 75^\circ$ rather than $\sim 85^\circ$ in [Fig. 10](#)). This may indicate that there is a further subdivision of the waves into different altitude regimes, possibly associated with secondary or tertiary AGWs created by the breaking of primary AGWs (e.g. [Vadas & Fritts, 2002](#); [Vadas et al., 2003](#)), meaning that no single altitude can adequately explain their behaviour. Alternatively, the remaining waves may still be plasmaspheric but simply not field-aligned, which would explain their existence at periods which are unlikely for AGWs in the thermosphere but would require an alternative generation mechanism to be proposed.

5 Conclusions

The calibration solutions derived as part of astronomical observations with LOFAR provide a new source of TID observations. Due to the extremely high sensitivity of the measurements, they reveal parts of the TID population which may be inaccessible with other more traditional techniques. However, this does not mean they are without limitations, such as the lack of altitude information and limited spatial and temporal coverage. This highlights the need for a range of measurements to fully characterise the TID population, and for data from any given technique to be interpreted carefully to avoid confusing a bias or limitation of the technique for a characteristic of the TIDs themselves.

Using the LOFAR calibration solutions and the analysis method developed by Boyde et al. (2024), it has been shown that TIDs and other wave-like disturbances are ubiquitous above LOFAR during the daytime. No relationship with geomagnetic or solar activity could be identified, but this must be qualified by acknowledging that the observations span a relatively limited range of geomagnetic activity levels and only part of a solar cycle. Overall, the diurnal and seasonal occurrence rates show a complicated pattern, with different period ranges showing activity peaks at different times. However, generally the high sensitivity of the measurements made with LOFAR indicates the variations in TID occurrence with time and season are less dramatic than has been suggested previously.

In terms of wave parameters, LOFAR reveals wave activity across a wide range of wavelengths, phase velocities and propagation directions. As far as the wavelengths are concerned, it is difficult to make definite statements about relative occurrence rates for waves of different wavelength due to the limitations on LOFAR's sensitivity imposed by the finite baseline lengths and relatively short observing windows as discussed in Section 3.2. The wave amplitudes on average behave exactly as would be expected, with peak amplitudes around midday or just afterwards, and increasing amplitude with increasing solar activity, due to the corresponding increases in ionisation. As a given AGW induces a given relative perturbation in electron density (Hooke, 1968), the increase in the background ionisation causes an increase in the observed absolute amplitudes.

The wave propagation directions show clear evidence that the dominant mechanism controlling them is filtering by thermospheric winds. Comparison across period ranges has indicated that the altitude at which the observations best match the climatological winds varies with period, with shorter periods corresponding to lower altitudes. As LOFAR does not provide any direct altitude information, it is unclear from this data whether the shorter period waves are actually being observed at lower altitudes or if they are simply more controlled by winds at lower altitudes due to having longer vertical wavelengths, as discussed in Section 4.2. For the periods below 10 min, this relationship with the neutral winds breaks down, as might be expected given that these waves would be below the Brunt-Väisälä period at the altitude of the F-region ionosphere and hence less likely to be explained by thermospheric AGWs. The possible contribution of plasmaspheric waves in this period range, especially below 5 min, as discussed in Section 4.3 may provide a partial explanation for the significant differences to behaviour at longer periods.

In general, the observations with LOFAR show that the dominant wave population detected are TIDs associated with upward propagating AGWs. There is no apparent relationship with geomagnetic activity, suggesting that the waves are primarily driven from within the Earth system. These observations therefore provide a means of investigating vertical coupling between the lower atmosphere and thermosphere.

Acknowledgements

This work is supported by the Leverhulme Trust under Research Project Grant RPG-2020-140. Ben Boyde acknowledges receipt of a PhD studentship from the same grant. This paper is based on data obtained with the International LOFAR Telescope (ILT). LOFAR (van Haarlem et al., 2013) is the Low Frequency Array designed and constructed by ASTRON. It has observing, data processing, and data storage facilities in several countries, that are owned by various parties (each with their own funding sources), and that are collectively operated by the ILT foundation under a joint scientific policy. The ILT resources have benefited from the following recent major funding sources: CNRS-INSU, Observatoire de Paris and Université d'Orléans, France; BMBF, MIWF-NRW, MPG, Germany; Science Foundation Ireland (SFI), Department of Business, Enterprise and Innovation (DBEI), Ireland; NWO, The Netherlands; The Science and Technology Facilities Council (STFC), UK; Ministry of Science and Higher Education, Poland. LOFAR data are available at <https://lta.lofar.eu/>. The ap data used here was obtained from GFZ Potsdam (<https://www.gfz-potsdam.de/en/section/geomagnetism/data-products-services/geomagnetic-kp-index>). The editor thanks Michael Kosch and an anonymous reviewer for their assistance in evaluating this paper.

References

- Alken P, Thébaud E, Beggan CD, Amit H, Aubert J, et al. 2021. International geomagnetic reference field: the thirteenth generation. *Earth Planets Space* **73**(1): 1–25. <https://doi.org/10.1186/s40623-020-01288-x>.
- Beser K, Mevius M, Grzesiak M, Rothkaehl H. 2022. Detection of periodic disturbances in LOFAR calibration solutions. *Remote Sens* **14**(7): 1719. <https://doi.org/10.3390/rs14071719>.
- Boyde B, Wood A, Dorrian G, Fallows RA, Themens D, et al. 2022. Lensing from small-scale travelling ionospheric disturbances observed using LOFAR. *J Space Weather Space Clim* **12**: 34. <https://doi.org/10.1051/swsc/2022030>.
- Boyde B, Wood A, Dorrian G, Sweijen F, de Gasperin F, Mevius M, Beser K, Themens D. 2024. Wavelet analysis of differential TEC measurements obtained using LOFAR. *Radio Sci* **59**(4): e2023RS007871. <https://doi.org/10.1029/2023RS007871>.
- Buss S, Hertzog A, Hostettler C, Bui TB, Lüthi D, Wernli H. 2004. Analysis of a jet stream induced gravity wave associated with an observed ice cloud over Greenland. *Atmos Chem Phys* **4**(5): 1183–1200. <https://doi.org/10.5194/acp-4-1183-2004>.
- Chum J, Šindelářová T, Laštovička J, Hruška F, Burešová D, Baše J. 2010. Horizontal velocities and propagation directions of gravity waves in the ionosphere over the Czech Republic. *J Geophys Res* **115**: A11322. <https://doi.org/10.1029/2010JA015821>.
- Chum J, Podolská K, Ruzs J, Baše J, Tedoradze N. 2021. Statistical investigation of gravity wave characteristics in the ionosphere. *Earth Planets Space* **73**: 1–16. <https://doi.org/10.1186/s40623-021-01379-3>.
- Cowling D, Webb H, Yeh K. 1971. Group rays of internal gravity waves in a wind-stratified atmosphere. *J Geophys Res* **76**(1): 213–220. <https://doi.org/10.1029/JA076i001p00213>.

- Crowley G, Jones T, Dudeney J. 1987. Comparison of short period TID morphologies in Antarctica during geomagnetically quiet and active intervals. *J Atmos Sol Terr Phys* **49(11)**: 1155–1162. [https://doi.org/10.1016/0021-9169\(87\)90098-5](https://doi.org/10.1016/0021-9169(87)90098-5).
- Crowley G, Rodrigues FS. 2012. Characteristics of traveling ionospheric disturbances observed by the TIDDBIT sounder. *Radio Sci* **47(4)**: RSOL22. <https://doi.org/10.1029/2011RS004959>.
- de Gasperin F, Dijkema TJ, Drabent A, Mevius M, Rafferty D, et al. 2019. Systematic effects in LOFAR data: A unified calibration strategy. *A&A* **622**: A5. <https://doi.org/10.1051/0004-6361/201833867>.
- de Gasperin F, Mevius M, Rafferty D, Intema H, Fallows R. 2018. The effect of the ionosphere on ultra-low-frequency radio-interferometric observations. *A&A* **615**: A179. <https://doi.org/10.1051/0004-6361/201833012>.
- de Gasperin F, Williams WL, Best P, Brüggem M, Brunetti G, et al. 2021. The LOFAR LBA Sky Survey-I. Survey description and preliminary data release. *A&A* **648**: A104. <https://doi.org/10.1051/0004-6361/202140316>.
- Ding F, Wan W, Xu G, Yu T, Yang G, Wang J-S. 2011. Climatology of medium-scale traveling ionospheric disturbances observed by a GPS network in central China. *J Geophys Res Space Phys* **116(A9)**: A09327. <https://doi.org/10.1029/2011JA016545>.
- Dorrian G, Fallows R, Wood A, Themens DR, Boyde B, Krankowski A, Bisi M, Dabrowski B, Vocks C. 2023. LOFAR observations of substructure within a traveling ionospheric disturbance at mid-latitude. *Space Weather* **21(1)**: e2022SW003198. <https://doi.org/10.1029/2022SW003198>.
- Drob DP, Emmert JT, Meriwether JW, Makela JW, Doornbos E, et al. 2015. An update to the Horizontal Wind Model (HWM): The quiet time thermosphere. *Earth Space Sci* **2(7)**: 301–319. <https://doi.org/10.1002/2014EA000089>.
- Ester M, Kriegel H-P, Sander J, Xu X. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In: Simoudis E, Han J, Fayyad U (Eds.), *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, KDD-96*, AAAI Press, Menlo Park, CA, USA, pp. 226–231. <https://www.aaai.org/Papers/KDD/1996/KDD96-037.pdf>.
- Fallows RA, Forte B, Astin I, Allbrook T, Arnold A, et al. 2020. A LOFAR observation of ionospheric scintillation from two simultaneous travelling ionospheric disturbances. *J Space Weather Space Clim* **10**: 10. <https://doi.org/10.1051/swsc/2020010>.
- Fedorenko YP, Fedorenko V, Lysenko V. 2011. Parameters of the medium-scale traveling ionospheric disturbances model deduced from measurements. *Geomagn Aeron* **51(1)**: 88–104. <https://doi.org/10.1134/S0016793210061015>.
- Fišer J, Chum J, Liu J-Y. 2017. Medium-scale traveling ionospheric disturbances over Taiwan observed with HF Doppler sounding. *Earth Planets Space* **69**: 1–10. <https://doi.org/10.1186/s40623-017-0719-y>.
- Frissell NA, Baker J, Ruohoniemi JM, Gerrard AJ, Miller ES, Marini JP, West ML, Bristow WA. 2014. Climatology of medium-scale traveling ionospheric disturbances observed by the midlatitude Blackstone SuperDARN radar. *J Geophys Res Space Phys* **119(9)**: 7679–7697. <https://doi.org/10.1002/2014JA019870>.
- Frissell NA, Baker JBH, Ruohoniemi JM, Greenwald RA, Gerrard AJ, Miller ES, West ML. 2016. Sources and characteristics of medium-scale traveling ionospheric disturbances observed by high-frequency radars in the North American sector. *J Geophys Res Space Phys* **121(4)**: 3722–3739. <https://doi.org/10.1002/2015JA022168>.
- Fritts DC, Alexander MJ. 2003. Gravity wave dynamics and effects in the middle atmosphere. *Rev Geophys* **41(1)**: 1003. <https://doi.org/10.1029/2001RG000106>.
- Galushko V, Paznukhov V, Yampolski Y, Foster J. 1998. Incoherent scatter radar observations of AGW/TID events generated by the moving solar terminator. *Ann Geophys* **16(7)**: 821–827. <https://doi.org/10.1007/s00585-998-0821-3>.
- Hines CO. 1960. Internal atmospheric gravity waves at ionospheric heights. *Can J Phys* **38(11)**: 1441–1481. <https://doi.org/10.1139/p60-150>.
- Hocke K, Schlegel K. 1996. A review of atmospheric gravity waves and travelling ionospheric disturbances: 1982–1995. *Ann Geophys* **14(9)**: 917. <https://doi.org/10.1007/s00585-996-0917-6>.
- Hoogeveen G, Jacobson A. 1997. Improved analysis of plasmasphere motion using the VLA radio interferometer. *Ann Geophys* **15(2)**: 236–245. <https://doi.org/10.1007/s00585-997-0236-6>.
- Hooke WH. 1968. Ionospheric irregularities produced by internal atmospheric gravity waves. *J Atmos Sol Terr Phys* **30(5)**: 795–823. [https://doi.org/10.1016/S0021-9169\(68\)80033-9](https://doi.org/10.1016/S0021-9169(68)80033-9).
- Hunsucker RD. 1982. Atmospheric gravity waves generated in the high-latitude ionosphere: a review. *Rev Geophys* **20(2)**: 293–315. <https://doi.org/10.1029/RG020i002p00293>.
- Inchin PA, Snively JB, Zettergren MD, Komjathy A, Verkhoglyadova OP, Tulasi Ram S. 2020. Modeling of ionospheric responses to atmospheric acoustic and gravity waves driven by the 2015 Nepal 7.8 Gorkha earthquake. *J Geophys Res Space Phys* **125(4)**: e2019JA027200. <https://doi.org/10.1029/2019JA027200>.
- Jacobson AR, Carlos RC, Massey RS, Wu G. 1995. Observations of traveling ionospheric disturbances with a satellite-beacon radio interferometer: seasonal and local time behavior. *J Geophys Res Space Phys* **100(A2)**: 1653–1665. <https://doi.org/10.1029/94JA02663>.
- Jacobson AR, Erickson WC. 1992. Wavenumber-resolved observations of ionospheric waves using the very large array radiotelescope. *Planet Space Sci* **40(4)**: 447–455. [https://doi.org/10.1016/0032-0633\(92\)90163-I](https://doi.org/10.1016/0032-0633(92)90163-I).
- Jacobson AR, Erickson WC. 1993. Observations of electron density irregularities in the plasmasphere using the VLA radio interferometer. *Ann Geophys* **11(10)**: 869–888.
- Kirkland MW, Jacobson AR. 1998. Drift-parallax determination of the altitude of traveling ionospheric disturbances observed with the Los Alamos radio-beacon interferometer. *Radio Sci* **33(6)**: 1807–1825. <https://doi.org/10.1029/98RS02033>.
- Koval A, Chen Y, Stanislavsky A, Zhang Q-H. 2017. Traveling ionospheric disturbances as huge natural lenses: Solar radio emission focusing effect. *J Geophys Res Space Phys* **122(9)**: 9092–9101. <https://doi.org/10.3847/1538-4357/ab1b52>.
- Lane TP, Reeder MJ, Clark TL. 2001. Numerical modeling of gravity wave generation by deep tropical convection. *J Atmos Sci* **58(10)**: 1249–1274. [https://doi.org/10.1175/1520-0469\(2001\)058<1249: NMOGWG>2.0.CO;2](https://doi.org/10.1175/1520-0469(2001)058<1249: NMOGWG>2.0.CO;2).
- Laughman B, Fritts DC, Lund TS. 2017. Tsunami-driven gravity waves in the presence of vertically varying background and tidal wind structures. *J Geophys Res Atmos* **122(10)**: 5076–5096. <https://doi.org/10.1002/2016JD025673>.
- Liu X, Xu J, Yue J, Vadas SL, Becker E. 2019. Orographic primary and secondary gravity waves in the middle atmosphere From 16-year SABER observations. *Geophys Res Lett* **46(8)**: 4512–4522. <https://doi.org/10.1029/2019GL082256>.
- Medvedev A, Ratovsky K, Tolstikov M, Alsatkin S, Shcherbakov A. 2015. A statistical study of internal gravity wave characteristics using the combined Irkutsk incoherent scatter radar and digisonde data. *J Atmos Sol Terr Phys* **132**: 13–21. <https://doi.org/10.1016/j.jastp.2015.06.012>.
- Medvedev AV, Ratovsky KG, Tolstikov MV, Oinats AV, Alsatkin SS, Zherebtsov GA. 2017. Relation of internal gravity wave anisotropy with neutral wind characteristics in the upper atmosphere. *J Geophys Res Space Phys* **122(7)**, 7567–7580. <https://doi.org/10.1002/2017JA024103>.

- Mevius M, van der Tol S, Pandey VN, Vedantham HK, Brentjens MA, et al. 2016. Probing ionospheric structures using the LOFAR radio telescope. *Radio Sci* **51**(7): 927–941. <https://doi.org/10.1002/2016RS006028>.
- Negrea C, Zobotin N, Bullett T. 2018. Seasonal variability of the midlatitude traveling ionospheric disturbances from Wallops Island, VA, dynasonde data: evidence of a semiannual variation. *J Geophys Res Space Phys* **123**(6), 5047–5054. <https://doi.org/10.1029/2017JA025164>.
- Negrea C, Zobotin N, Bullett T, Fuller-Rowell T, Fang T-W, Codrescu M. 2016. Characteristics of acoustic gravity waves obtained from dynasonde data. *J Geophys Res Space Phys* **121**(4): 3665–3680. <https://doi.org/10.1002/2016JA022495>.
- Nicolls MJ, Heinselman CJ. 2007. Three-dimensional measurements of traveling ionospheric disturbances with the Poker Flat incoherent scatter radar. *Geophys Res Lett* **34**(21): L21104. <https://doi.org/10.1029/2007GL031506>.
- Oinats AV, Nishitani N, Ponomarenko P, Berngardt OI, Ratovsky KG. 2016. Statistical characteristics of medium-scale traveling ionospheric disturbances revealed from the Hokkaido East and Ekaterinburg HF radar data. *Earth Planets Space* **68**(1): 1–13. <https://doi.org/10.1186/s40623-016-0390-8>.
- Oliver WL, Otsuka Y, Sato M, Takami T, Fukao S. 1997. A climatology of F region gravity wave propagation over the middle and upper atmosphere radar. *J Geophys Res Space Phys* **102**(A7): 14499–14512. <https://doi.org/10.1029/97JA00491>.
- Otsuka Y, Suzuki K, Nakagawa S, Nishioka M, Shiokawa K, Tsugawa T. 2013. GPS observations of medium-scale traveling ionospheric disturbances over Europe. *Ann Geophys* **31**(2), 163–172. <https://doi.org/10.5194/angeo-31-163-2013>.
- Panasenko SV, Goncharenko LP, Erickson PJ, Aksonova KD, Domin IF. 2018. Traveling ionospheric disturbances observed by Kharkiv and Millstone Hill incoherent scatter radars near vernal equinox and summer solstice. *J Atmos Sol Terr Phys* **172**: 10–23. <https://doi.org/10.1016/j.jastp.2018.03.001>.
- Park J, Lühr H, Lee C, Kim YH, Jee G, Kim J-H. 2014. A climatology of medium-scale gravity wave activity in the midlatitude/low-latitude daytime upper thermosphere as observed by CHAMP. *J Geophys Res Space Phys* **119**: 2187–2196. <https://doi.org/10.1002/jgra.50886>.
- Sato K, Yoshiki M. 2008. Gravity wave generation around the polar vortex in the stratosphere revealed by 3-hourly radiosonde observations at Syowa Station. *J Atmos Sci* **65**(12): 3719–3735. <https://doi.org/10.1175/2008JAS2539.1>.
- Sivakandan M, Otsuka Y, Ghosh P, Shinagawa H, Shinbori A, Miyoshi Y. 2021. Comparison of seasonal and longitudinal variation of daytime MSTID activity using GPS observation and GAIA simulations. *Earth Planets Space* **73**: 1–16. <https://doi.org/10.1186/s40623-021-01369-5>.
- Snively JB, Pasko VP. 2003. Breaking of thunderstorm-generated gravity waves as a source of short-period ducted waves at mesopause altitudes. *Geophys Res Lett* **30**(24): 2254. <https://doi.org/10.1029/2003GL018436>.
- Thompson AR, Clark B, Wade C, Napier PJ. 1980. The very large array. *Astrophys J Suppl Ser* **44**: 151–167. <https://doi.org/10.1086/190688>.
- Torrence C, Compo GP. 1998. A practical guide to wavelet analysis. *Bull Am Meteorol Soc* **79**(1): 61–78. [https://doi.org/10.1175/1520-0477\(1998\)079<0061:APGTWA>2.0.CO;2](https://doi.org/10.1175/1520-0477(1998)079<0061:APGTWA>2.0.CO;2).
- Trigg H, Dorrian G, Boyde B, Wood A, Fallows RA, Mevius M. 2024. Observations of high definition symmetric quasi-periodic scintillations in the mid-latitude ionosphere with LOFAR. *J Geophys Res Space Phys* **129**(7): e2023JA032336. <https://doi.org/10.1029/2023JA032336>.
- Vadas SL. 2007. Horizontal and vertical propagation and dissipation of gravity waves in the thermosphere from lower atmospheric and thermospheric sources. *J Geophys Res Space Phys* **112**(A6): A06305. <https://doi.org/10.1029/2006JA011845>.
- Vadas SL, Fritts DC. 2002. The importance of spatial variability in the generation of secondary gravity waves from local body forces. *Geophys Res Lett* **29**(20): 45-1–45-4. <https://doi.org/10.1029/2002GL015574>.
- Vadas SL, Fritts DC. 2005. Thermospheric responses to gravity waves: Influences of increasing viscosity and thermal diffusivity. *J Geophys Res Atmos* **110**(D15): D15103. <https://doi.org/10.1029/2004JD005574>.
- Vadas SL, Fritts DC, Alexander MJ. 2003. Mechanism for the generation of secondary waves in wave breaking regions. *J Atmos Sci* **60**(1): 194–214. [https://doi.org/10.1175/1520-0469\(2003\)060<0194:MFTGOS>2.0.CO;2](https://doi.org/10.1175/1520-0469(2003)060<0194:MFTGOS>2.0.CO;2).
- van Haarlem MP, Wise MW, Gunst AW, Heald G, McKean JP, et al. 2013. LOFAR: The LOw-Frequency ARray. *A&A* **556**: A2. <https://doi.org/10.1051/0004-6361/201220873>.
- Waldock J, Jones T. 1986. HF Doppler observations of medium-scale travelling ionospheric disturbances at mid-latitudes. *J Atmos Sol Terr Phys* **48**(3): 245–260. [https://doi.org/10.1016/0021-9169\(86\)90099-1](https://doi.org/10.1016/0021-9169(86)90099-1).
- Wood AG, Dorrian GD, Boyde B, Fallows RA, Themens DR, Mevius M, Sprenger T, Main R, Pryse SE, Elvidge S. 2024. Quasi-stationary substructure within a sporadic E layer observed by the Low-Frequency Array (LOFAR). *J Space Weather Space Clim* **14**: 27. <https://doi.org/10.1051/swsc/2024024>.
- Wright CJ, Hindley NP, Alexander MJ, Barlow M, Hoffmann L, et al. 2022. Surface-to-space atmospheric waves from Hunga Tonga–Hunga Ha’apai eruption. *Nature* **609**(7928): 741–746. <https://doi.org/10.1038/s41586-022-05012-5>.
- Wüst S, Bittner M. 2006. Non-linear resonant wave-wave interaction (triad): Case studies based on rocket data and first application to satellite data. *J Atmos Sol Terr Phys* **68**(9): 959–976. <https://doi.org/10.1016/j.jastp.2005.11.011>.
- Xu S, Vadas SL, Yue J. 2024. Quiet time thermospheric gravity waves observed by GOCE and CHAMP. *J Geophys Res Space Phys* **129**(1): e2023JA032078. <https://doi.org/10.1029/2023JA032078>.
- Yiğit E, Knížová PK, Georgieva K, Ward W. 2016. A review of vertical coupling in the Atmosphere-ionosphere system: effects of waves, sudden stratospheric warmings, space weather, and of solar activity. *J Atmos Sol Terr Phys* **141**: 1–12. <https://doi.org/10.1016/j.jastp.2016.02.011>.
- Yin F, Lühr H, Park J, Wang L. 2019. Comprehensive analysis of the magnetic signatures of small-scale traveling ionospheric disturbances, as observed by Swarm. *J Geophys Res Space Phys* **124**(12): 10794–10815. <https://doi.org/10.1029/2019JA027523>.
- Zalizovski AV, Yampolski YM, Mishin E, Kashcheyev SB, Sopin AO, Koloskov AV, Lisachenko VN, Reznichenko AI. 2021. Multi-position facility for HF Doppler sounding of ionospheric inhomogeneities in Ukraine. *Radio Sci* **56**(10): e2021RS007303. <https://doi.org/10.1029/2021RS007303>.

Cite this article as: Boyde B, Wood A, Dorrian G, de Gasperin F & Mevius M, et al. 2025. Statistics of travelling ionospheric disturbances observed using the LOFAR radio telescope. *J. Space Weather Space Clim.* **15**, 6. <https://doi.org/10.1051/swsc/2025002>.

Appendix

Magnetic field shadows

To calculate the orientation of the “magnetic field shadow” (the apparent orientation of field aligned structures when projected into the LOFAR ground plane) for a given observation, first the co-ordinates of the line of sight at a given altitude were calculated following the method given by Dorrian et al. (2023). Once the geographic latitude and longitude corresponding to the specified altitude were known, the magnetic field orientation (inclination and declination) at that location were obtained from the IGRF. This provided the magnetic field vector in the local Cartesian co-ordinates (East, North and vertical). In order to compare this to the line of sight direction, defined in terms of source azimuth and elevation as viewed from LOFAR, the magnetic field vector must therefore be converted into the local Cartesian co-ordinates of LOFAR. For the purposes of these calculations, the location of LOFAR was taken to be that of station CS002, the most central of the core stations.

The conversion from the local Cartesian co-ordinates in which the magnetic field is defined (at a latitude and longitude of λ_i, ϕ_i) and those of LOFAR (latitude and longitude of λ_0, ϕ_0) is non-trivial to carry out directly. For simplicity, we defined a universal co-ordinate system in which the z co-ordinate points along the Earth’s rotation axis towards the North pole, the x co-ordinate points through 0° longitude and the y co-ordinate points through 90° E longitude. The transformation to and from any local Cartesian system to this universal system was then defined allowing transformations between any two local Cartesian systems in a two step process. The co-ordinate vectors at (λ_i, ϕ_i) were defined as \hat{E}' , \hat{N}' and \hat{Z}' respectively, and those at (λ_0, ϕ_0) as \hat{E} , \hat{N} and \hat{Z} respectively. The co-ordinate vectors in the universal co-ordinate system were \hat{x} , \hat{y} and \hat{z} respectively.

The transformation from (λ_i, ϕ_i) to the universal co-ordinate system can then be expressed as:

$$\begin{aligned}\hat{x} &= -\sin(\phi_i)\hat{E}' + \cos(\phi_i)\left[\cos(\lambda_i)\hat{Z}' - \sin(\lambda_i)\hat{N}'\right], \\ \hat{y} &= \cos(\phi_i)\hat{E}' + \sin(\phi_i)\left[\cos(\lambda_i)\hat{Z}' - \sin(\lambda_i)\hat{N}'\right], \\ \hat{z} &= \cos(\lambda_i)\hat{N}' + \sin(\lambda_i)\hat{Z}'.\end{aligned}\quad (\text{A1})$$

The transformation from the universal co-ordinate system to (λ_0, ϕ_0) is then:

$$\begin{aligned}\hat{E} &= \cos(\phi_0)\hat{y} - \sin(\phi_0)\hat{x}, \\ \hat{N} &= \cos(\lambda_0)\hat{z} - \sin(\lambda_0)\left[\cos(\phi_0)\hat{x} + \sin(\phi_0)\hat{y}\right], \\ \hat{Z} &= \sin(\lambda_0)\hat{z} + \cos(\lambda_0)\left[\cos(\phi_0)\hat{x} + \sin(\phi_0)\hat{y}\right].\end{aligned}\quad (\text{A2})$$

Once the magnetic field vector was transformed into the co-ordinate system at (λ_0, ϕ_0) its “shadow” was then calculated given the orientation of the line of sight. This “shadow” is the vector purely in the horizontal (E-N) plane which has the same projection onto the plane perpendicular to the line of sight as the magnetic field vector itself. To calculate this the magnetic field vector \vec{B} was added to a linear multiple of the line of sight vector \vec{l} such that the resulting vector \vec{B}' had no component in \hat{Z} . This can be expressed mathematically as

$$\vec{B}' = \vec{B} - \frac{B_Z}{l_Z}\vec{l}, \quad (\text{A3})$$

where B_Z and l_Z are the components along \hat{Z} of \vec{B} and \vec{l} respectively.

To identify potential field-aligned structures, this calculation was initially carried out for altitudes from $0.1R_E$ to $10.1R_E$ in steps of $1R_E$ (R_E denotes the radius of the Earth) for each wave event. From these, either the wave was classified as non-field-aligned (no altitudes for which \vec{B}' was close to perpendicular to the wave propagation) or Newton-Raphson iteration was used to determine the altitude at which the structure would align with the field (using a tenth of the estimated azimuth uncertainty of the wave as a convergence criterion, the Monte Carlo randomly generated azimuths used a tenth of the mean of the estimated azimuth uncertainties which was approximately 0.5°).