

Reply to Reviewer 2

We are very grateful to the Reviewer for the evaluation of this manuscript and for raising very interesting points and suggestions.

Before implementing any changes, we have prepared some answers and reflections, together with suggestions of possible changes that could be made.

We wonder whether these may be in the right direction?

The text of the Referee's report (in normal font type) is followed by our reply in Italic.

Review of egusphere-2026-1741

RC2: 'Comment on egusphere-2026-1741', Anonymous Referee #2, 22 Aug 2026

Review of Forte et al

“First demonstration of the combination of observations and datasets from diverse instruments (in a coordinated experiment using KAIRA, EISCAT, GNSS) that enables the extraction of new information about the ionosphere”

10.5194/egusphere-2026-1741

Summary:

The paper describes analysis of a multi-sensor experiment combining radio astronomy telescope data (KAIRA in Finland), incoherent scatter radar altitude-resolved electron density profiles from both mainland UHF EISCAT and EISCAT Svalbard Radar locations, ionosonde data, and multiple-channel GNSS total electron content data. It focuses on spatial and temporal ionospheric remote sensing in both phase and amplitude variations / scintillations using coordinated experiments spread over several days. Indirect references are also made to observations from the SuperDARN HF radar network which determine the auroral convection boundaries for context.

This is a quite nice and detailed analysis that uses some sensor fusion concepts to determine spatial and temporal ionospheric variations. It is very good to see diagnostics from both the radio astronomy and ionospheric community being looked at simultaneously; this is the future of ionospheric remote sensing approaches in the modern era. KAIRA has been used before successfully as an opportunistic multi static large aperture receiver deriving ionospheric radar profiles with EISCAT as the transmitter. This study employs the KAIRA sensitive receivers and phased arrays by examining CasA apparent variations to detect fluctuation thresholds that are smaller than the resolution cells and noise floors offered by either the ISR or GNSS techniques. Various instruments are effectively shown to detect different types of ionospheric features.

I think the analysis is quite solid and comprehensive, and it will be a good demonstration for the community of the need for such fusion type analyses across multiple fields (astronomy, ionosphere). There are some issues that should be examined, and these are detailed below.

Many thanks for your comments and suggestions.

Detailed comments:

The title is way too long. I recommend something like “Multi-sensor fusion of radio astronomy and active ionospheric diagnostics for advanced spatial and temporal resolution in ionospheric remote sensing.”

Yes, it is rather long. Within the reply to Reviewer 1, we have proposed a new title as follows:

“Combining observations from different instruments (KAIRA all-sky imaging, EISCAT, GNSS) for a novel ionospheric sounding.”

However, the suggestion by Reviewer 2 seems quite interesting. Hence, a new title could be a combination of these two proposals, for example:

“Multi-sensor fusion of radio astronomy and traditional ionospheric diagnostics for novel spatial and temporal resolution in ionospheric remote sensing.”

Line 96 “Whilst phase fluctuations can be generated..” It would be good to insert a reference directly here (I think it is done later) on why the amplitude fluctuations only occur on small spatial scales as opposed to the phase fluctuation in the ‘frozen in’ approximation.

We propose to add a reference to Yeh and Liu (1982), where these aspects were discussed and reviewed in great detail.

Line 118 “This is also true for proposed systems on the far side of the Moon (Gorgolewski, 1965),..” I am not familiar with the 1965 reference, but I think the lunar analogy is reasonably weak. The text correctly states that the plasma frequency in the lunar environment is much lower - after all, several missions or mission concepts are aiming for radio astronomy at < 10 MHz from the lunar surface - but is this argument really needed here to strengthen the relevance case? I think it may not be needed; something to consider.

We agree that this argument is not strictly necessary to strengthen the case for this particular study. It is offered to provide context on future instruments, which will operate at a much higher sensitivity (e.g., SKAO) and in environments that are not yet fully understood. If not too problematic, we would be inclined to keep this paragraph.

Line 131 “On the other hand, higher scintillation” do you mean stronger scintillation?

Changed to “stronger”.

Line 157 “and even to observations in the middle latitude ionosphere in the presence of ionospheric irregularities. “. Why just mid-latitudes? What is special about them as opposed to tropical ionospheric locations near the equatorial ionospheric

anomaly, or even equatorial irregularities driven by the equatorial electrojet? I am having a hard time understanding the particular mid-latitude qualifier.

This sentence was indeed not very clear. The reference was to satellite signals at radio wave frequencies lower than L band. Historically, these would experience scintillation at middle latitudes too. However, this part is not essential and we propose to rephrase that sentence as:

“The methodology outlined here can indeed be applied to any LOFAR imaging and LOFAR beamformed VHF scintillation observations as well as to similar observations from other radio telescopes.”

Line 240 "Here, an interval of 3 minutes was utilised under the assumption of ergodicity". I realize you have given two references immediately following this sentence as to why the 3 minute integration time was used for S4, but rather than make the reader go find the reason, I recommend a few words here to quickly justify this choice, as it does underpin a lot of the following analysis.

We propose to add the following text after this sentence:

“...: this choice was used to capture the contribution from the spatial scales in the inertial subrange (which induce intensity scintillation) and their temporal variability in the presence of typical drift velocities in the ionosphere.”

Line 244: “the scintillation timescale appeared to be somewhat longer, hence introducing larger uncertainties in the S4 estimate.” This brings up a problem for the reader: you explicitly reference uncertainty values on S4 here, but none are given in the paper. Can you provide some kind of uncertainty estimate; otherwise, it’s hard to be more than qualitative about the results because I can’t ascertain what is statistically significant and what is instrumental uncertainty.

This sentence was indeed not very clear. We propose to delete it given the explanation provided with the previous sentence. A discussion of uncertainties in S4 would depend on the detailed knowledge of the spatial scales originating scintillation along a given ray path together with detailed knowledge of their statistical properties, making such a discussion not possible. For this reason, the S4 index is typically estimated by means of a standard deviation over a given interval of time under the assumption of ergodicity, given that it is not possible to calculate it over the ensemble. However, the process is not ergodic and the time interval needs to be capable of capturing the scintillation arising from the irregularities with spatial scales in the VHF inertial subrange as well as their temporal variability. As demonstrated previously (Flisek et al., 2023), the choice of 3 minutes appears to satisfy this requirement in the content of estimates of the S4 scintillation index from beamformed radio-telescope scintillation observations.

Line 249: “KAIRA recorded cross-correlation statistics with a time resolution of 1s at the frequency of 70.3 MHz “ Please provide rationale for why KAIRA chose that frequency for CasA monitoring since it of course is capable of using other frequencies.

We propose to add the following text in line 249:

“...: this value was chosen to maximise the resolution for the all-sky images from the LBA and, at the same time, to minimise the presence of absorption in the data.”

Line 283: “Since this flux is missing from the Low-frequency Sky Survey, which was produced carefully by measuring the zero-spacing flux and by controlling for the effects of mutual coupling between dipoles,” I am a bit confused here. Is there a reason why KAIRA did not remove the zero-spacing flux or include mutual coupling effects? Recommend some rewording to make it more clear that you are focusing on the region around CasA (as you clearly show later) and that these discrepancies are not relevant for the study presented in the paper, so it’s OK to have them.

KAIRA did not account for zero-spacing flux or mutual coupling effects because to do so requires a great deal of careful modelling.

We propose to add the following text at the end of line 283:

“However, this study is only concerned with the area surrounding the bright source CasA, and so these effects will not affect our ionospheric analysis.”

Line 320: “.. rotated from the source-centred coordinates by an angle θ (Fig. 4) to account for any rotation in the source image (the angle θ was estimated within the fit).” Theta measures the rotation ‘twist’ of the CasA image but it’s not clear at this point whether that angle data was presented later or physically interpreted. Perhaps make it more clear here: is the inclusion of that parameter meant to correct for celestial rotation effects relative to the principal axes of the KAIRA beam (e.g. astronomical)? Or is it a data stream that contains ionospheric information itself and how will it be interpreted?

The angle θ was not utilised in the analysis here. This angle was part of the preparation of the source image before applying the 2D gaussian fit. This is necessary to account for any rotation in the source image. As we are mainly interested in the displacement parameter, there was no attempt here to interpret any information carried by this angle. However, in a subsequent work that is being submitted, the rotation in the source image is utilised to relate source changes to ionospheric structures.

We propose to add the following text to that line:

“...: the angle θ was not utilised in the analysis afterwards.”

Line 336: “was fitted to the image area by applying a non-linear least squares method”. There is no discussion here of what happens when you encounter a CasA image that is warped by RFI impacts or other systematics. How do you reject this kind of data - χ^2 examination? Something else? Describing that here in a quick sentence or two will further give the reader confidence in your analyzed data. This would also help in the detailed explanation of the parameter fit that follows in the next 2-3 paragraphs.

The source-zoomed image is based on visibilities already normalised by the array median, before applying the 2D gaussian fit. This removes instrumental effects, including RFI.

We propose to add the modify line 293 as follows:


“The intensity $I(l, m)$ in all-sky images was normalised by the array-wise median (which removes instrumental effects, including RFI); then, a...”

Line 370 “In order to encapsulate the variations of the source image a new function that combines the temporal behaviour of the gaussian fit parameters is proposed”. Again as in an earlier comment, no quantitative uncertainty is given in these parameters (through e.g. error propagation). Can you give the reader any idea at all on the uncertainty of these coordinate-transformed parameters? I feel it’s important to gauge the numerical quality of the results.

Figure 7: Uncertainties?

This aspect (both at line 270 and within Figure 7) relates to the parameters deduced from the 2D gaussian fit and it appears to be best placed when these are introduced (right before line 370). We have estimated the uncertainties on these parameters in two ways: (1) by estimating the covariance in the (l, m) coordinates associated with the peak of the 2D gaussian fit and (2) by estimating the propagation of uncertainties from instrumental considerations.

Method 1 shows that the uncertainties are very small, of the order of approximately 10^{-2} degrees:



We could add these figures to the appendices, if helpful.

Method 2 is based on the propagation of uncertainties. McKay et al. (2015) give the SEFD of the array at 70 MHz as approximately 20kJy (their figure 11). This yields an RMS noise in the image of 32 Jy (e.g. Morgan & Ekers, 2021, Eq. (1)), compared with a brightness of Cas A at 74 MHz of 17694 Jy (Kassim et al. 2007).

Condon (1997), through Eq. (16), gives the errors on all fitted parameters for an elliptical Gaussian in terms of ρ , the product of signal-to-noise ratio (for this case, $17694/32 \approx 500$), and N , the number of independent samples in the fit.

For a source whose apparent size is comparable to the resolution of the instrument, N is approximately 1, and the fractional error on all fitted parameters is exceedingly small, again of the order of approximately 10^{-2} degrees (for D , the displacement, the fractional error is relative to the source size parametrised by σ_x and σ_y).

Consequently, even the smallest variability visible by eye in Figure 7 are large compared to the error induced by the noise; however the shortest timescale fluctuations we ascribe to instrumental variations.

We propose to add the following text before Line 370:

“The fractional error on the 2D gaussian fit parameters is exceedingly small, of the order of approximately 10^{-2} degrees. Considering that the RMS noise in the KAIRA all-sky image is approximately of the order of 32 Jy at 70 MHz (McKay et al., 2015; Morgan and Ekers, 2021) and considering the CasA brightness of approximately of the order of 17694 Jy at 74 MHz (Kassim et al., 2007), then the signal-to-noise ratio is approximately of the order of $17694/32 \approx 500$. The errors on all fitted parameters for an elliptical Gaussian are described in terms of ρ , the product of signal-to-noise (for this case, ~ 500), and N , the number of independent samples in the fit. For a source whose apparent size is comparable to the resolution of the instrument, N is approximately 1, and the fractional error on all fitted parameters is exceedingly small, of the order of approximately 10^{-2} degrees (for the displacement D , the fractional error is relative to the source size parametrised by σ_x and σ_y). Consequently, even the smallest variability visible by eye (e.g., later in Fig. 7) are large compared to the error induced by the noise; however, the shortest timescale fluctuations we ascribe to instrumental variations.”

With the following references:

*Morgan J. S. and Ekers R. (2021), A measurement of source noise at low frequency: Implications for modern interferometers. Publications of the Astronomical Society of Australia, Vol. 38, e013, pp. 1–6.
<https://doi.org/10.1017/pasa.2021.3>.*

Kassim N. E., T. Joseph, W. Lazio, W. C. Erickson, R. A. Perley, W. D. Cotton, E. W. Greisen, A. S. Cohen, B. Hicks, H. R. Schmitt, and D. Katz (2007), The 74 MHz System on the Very Large Array, The Astrophysical Journal Supplement Series, Vol. 172, N. 2, pp. 686-719, DOI 10.1086/519022.

Condon J. J. (1997), Errors in elliptical gaussian fits, Publications of the Astronomical Society of the Pacific, Vol. 109, N. 732, p. 166, DOI 10.1086/133871.

Line 410: “electron density (N_e) profiles were measured by means of the EISCAT UHF ISR in Tromsø “. Here is a big confusion for me. ISRs are the most comprehensive ground-based ionospheric diagnostic available due to their full altitude-resolved profiles. I can understand not using T_e and T_i information - but the ISR produces line-of-sight plasma velocity, resolved along the beam as a primary output. Why did this parameter not get used at all? Below, there are some comments relating to rather vague inferences on plasma motion from e.g. SuperDARN - but this precise measurement from ESR and UHF mainland wasn't used? It should be fully consistent with the other diagnostics. I strongly recommend including it. For instance, it has a projection along line-of-sight of the ionospheric convection velocity vector derived e.g. from SuperDARN and these should be consistent (and this will add further confidence to the multi-sensor results, especially when discussing the frozen-in phase screen approximation).

Line 440: “The Tromsø dynasonde results appeared to suggest (not shown here) the possible presence of a horizontal plasma drift, approximately between 100 – 300 m s⁻¹ during the observations, although this estimate is at a very low confidence.” See

above comment: why not seriously strengthen your arguments about plasma drift by using UHF ISR results?

In relation to both these points, the ion drift velocity estimated along the radar's beam was not utilised as this would relate to a component of the drift nearly orthogonal (under the geometry utilised here) to the horizontal drift, hence not relevant to the temporal scale observed, for example, in the measurements of VHF scintillation or of GNSS ROT.

Unfortunately, it was not possible to carry out tri-static measurements in this experiment (which would have enabled an estimate of the horizontal plasma drift).

We have analysed the ion drift velocity profiles and this velocity appears to be almost always of the order of approximately $100 \text{ m} \cdot \text{s}^{-1}$ on average throughout the observations, hence not providing any additional insight, unfortunately.

In the future, EISCAT_3D could provide more clarity on these aspects though it is frustrating that we do not seem to deduce an independent estimate from the EISCAT observations at this time.

Line 461: "making ROT a useful detector of phase fluctuations induced by ionospheric irregularities". Recommend adding "~23 cm" before the word "phase".

Added "L-band" before the word "phase".

Line 519: "The displacement D appears to show values smaller than the overall resolution in all-sky images achievable with KAIRA: this is due to the fact that the source-zoomed image was further discretised and pixelated. " A bit vague - can you quantify any further the lower limit on D (e.g. it is a minimum value) imposed by the known instrumental distortion / pixelation?

Due to the uncertainty analysis added before line 370 (as mentioned above), this sentence can be deleted.

Instead, we propose to add the following text at the end of line 525:

"...: the case from 24 November 2021 can be considered as a quiet reference case."

Line 564: "Throughout the observations (Fig. 1), EISCAT detected enhancements in the electron density in both the F region and the E region." I think it is quite important here (and important for the later conclusions) to state quantitatively what the along-beam range resolution is, which is known considering the transmit pulse length/code baud length, and the cross-beam spatial resolution, which is a factor of the beam FWHM angular value multiplied by the range. This is because you are combining radio astronomers with ionospheric scientists as readers, and the former do not know the typical values used by the latter in these types of experiments.

Based on the experiments utilised in this case (beata for ESR and bella for EISCAT), the steps in altitude as output by the radars are as follows.

Altitude values as present in “beata” data (an example, small differences exist between files):

78.18049 81.89321 85.60597 89.31877 93.03162 96.74450 100.45742
104.17038 107.88338 114.02432 119.21811 125.41738 132.94956 141.66769
150.72142 160.92503 173.23455 186.52292 199.27952 213.53922 228.69170
245.62998 263.53147 282.04962 301.37066 321.73209 343.37209 365.38519
388.23822 413.61869 439.87569 465.25172 492.39594 521.55628 551.28469
581.73438 612.85500 634.16381

Altitude values as present in “bella” data (an example, small differences exist between files):

77.08484 79.29824 82.60092 85.93859 88.15209 91.44424
94.79273 97.00633 100.32770 103.64726 105.86095 109.72074
114.87525 120.18995 126.27466 134.12503 142.83258 152.55703
163.19917 174.32072 186.61836 199.84563 214.50686 230.11611
246.51663 263.84731 282.81497 302.61187 323.11066 344.27944
366.29103 389.72191 414.13128 439.67231 466.42025 493.58869
522.55075 552.30706 582.87987 613.98063 646.01575 679.78937
714.34237 749.84475 786.29156 823.90069 862.68031 902.57181
942.64544 983.59713 1026.58825 1070.12200 1114.42863 1159.59150
1206.51150 1254.74212 1302.57575 1351.84488

We propose to add the following text at the end of line 570:

“The altitude resolution is nonlinear and varies between about 3 km at 100 km of altitude up to about 30 km at 500 km; an indication of the discretisation can be seen in the pixelization of the data in Fig. 9.”

Line 574: “the EISCAT/ESR electron density profiles were integrated between approximately 100 km and 400 km of altitude (Jakowski et al., 1996; Forte et al., 2013; Forte et al., 2017), thus providing an estimate of the slant TEC along the radars’ beams.” I am glad to see this quantity being calculated, but think it is a large mistake to name it TEC - it instantly confuses it with the much more widely known TEC coming from GNSS, which of course goes out to 20,000+ km. I recommend switching this name to “integrated electron content” or IEC to remove that confusion. (Note that “radar ROT” doesn’t have that confusion.)

Changed to “integrated electron content” and to “IEC”.

Figure 10: you really need error bars on this or at least a statement of the quantitative uncertainty. You can calculate it directly because the ISR provides uncertainties on each electron density measurement as a function of altitude. It helps with your descriptions to guide the reader to interpret the fluctuations as geophysical, not instrumental.

We have indeed analysed the uncertainties provided with the electron density profiles. The relative errors on the electron density values within the interval of altitudes that we are interested in are typically less than 1%.

We propose to add the following text after line 570:

“The relative standard deviation (RSD) of N_e values are typically much less than 1%”.

By propagating the uncertainties from the electron density values to the integrated electron density content, the relative standard deviation in the ROT values appear to be less than 1%, as well.

We propose to add the following text after line 581:

“The uncertainties in the ROT values are simply the RMS combination of those in N_e when integrated along the radar line-of-sight through the data. The result is that the RSD in ROT values are typically also less than 1%; thus, essentially negligible on the scale of the temporal variation of ROT itself.”

Line 609: “Figure 11 also shows the probability that |ROT| exceeds a given threshold”. This is really odd to me - you describe that “calculated by dividing the number of all the GNSS observations showing |ROT| greater than the threshold by the number of all available GNSS 30 s observations in any 3-minute interval.” That’s a percentage, not a probability, of ‘spikes’ at 30 second integration within the 3-minute block. Probability implies prediction to me and I don’t think that is what you are trying to convey.

Changed to “percentage”.

Also, the threshold of 0.2 TECU min⁻¹ - I got lost trying to figure out why that particular value and not 0.1, 0.3, or some other number.

This threshold was identified qualitatively on the basis of the GNSS observations. This value represent approximately the order of magnitude of the instrumental noise floor in this type of observations. Any gradient below this limit cannot be resolved by individual GNSS receivers. In response to a suggestion from Referee 1, this value has been also converted in units of TECU · km⁻¹, i.e. an indication of spatial gradients that can be resolved by individual GNSS receivers.

We propose to add the following text at the end of line 609:

“..., approximately the order of magnitude of the GNSS noise floor.”

Line 637: “When compared with EISCAT/ESR electron density profiles, enhancements in ROT seemed to be caused by structured ionisation extending over a wider interval of altitudes rather than by ionisation gradients occurring across thinner layers”. Can you reword this? I got confused trying to parse it out to

understand why clear ROT structures didn't show up at VHF, which is presumably more sensitive as you describe later.

Intensity scintillation is caused by small-scale irregularities drifting across the beam. The absence of VHF scintillation in the presence of GNSS ROT can be explained with the presence of ionisation structures that are elongated in altitude and where small-scale irregularities in the VHF inertial subrange are absent (the ionisation structures are smoother).

The following re-wording is proposed:

“Ionisation structures extending over a wider interval of altitudes seemed to favour enhancements in GNSS ROT. Thinner or smaller structures may not be capable of enhancing GNSS ROT (as in....”.

Line 641: “Whilst a horizontal plasma drift appeared to be plausible given the presence of intensity fluctuations on KAIRA scintillation observations (low-confidence observations from the Tromsø dynasonde appeared to support this possibility), GNSS ROT observations did not seem to show a clear evidence for that.” See earlier comment about ISR data on velocities. This also applies to Line 726 description of drifts.

As mentioned earlier, the ion drift velocity estimated along the radar's beam is along a direction nearly orthogonal (under the geometry utilised here) to the horizontal drift. For this reason, we cannot seem to be able to utilise the EISCAT ion drift profiles in a conclusive way, unfortunately.

Line 695: “Considering that the spatial resolution of EISCAT electron density profiles is approximately of the order of tens of kilometres at F region altitudes,..” See earlier comment about ISR beam dimensions: you can be much more quantitative than this here.

We propose to modify that sentence as:

“Considering that the spatial resolution of EISCAT electron density profiles is approximately of the order of 30 km at F region altitudes,...”.

Line 701: “Interestingly, a sporadic E layer could be observed to form from about 20:00 UTC on 25 November 2021” SpE layers are typically 1-2 km thick, but the radar shows them extending over 30-40 km. This must be an along-beam waveform resolution effect, and to avoid people discounting that identification, I would make a clear statement here about radar waveform resolution.

We propose to add the following text in that line:

“...:the resolution in this case could be limited, for example, by along-beam waveform resolution effects).”

Citation:

<https://eur01.safelinks.protection.outlook.com/?url=https%3A%2F%2Fdoi.org%2F10.>

5194%2Fegusphere-2026-1741-

RC2&data=05%7C02%7Cbf263%40bath.ac.uk%7C5f1b83c6e9a4411df05a08df01c2
c3e6%7C377e3d224ea1422db0ad8fcc89406b9e%7C0%7C0%7C63923160815283
0108%7CUnknown%7CTWFpbGZsb3d8eyJFbXB0eU1hcGkiOnRydWUsIIYiOilwLjA
uMDAwMCIsIIAiOiJXaW4zMilslkFOljoiTWFpbClslldUljoyfQ%3D%3D%7C0%7C%7
C%7C&sdata=Ptpi2Cnt2AqKX4TehpALvsLfhS%2BYuRATE4ZJGeNhWjY%3D&rese
rved=0