

Manuscript egusphere-2025-5374: Dune aurora: Statistical survey from a citizen science database

Response to Reviewer #1

We thank the Reviewer for their thorough assessment of our manuscript and for raising important points to improve the paper. We truly appreciate the time and effort spent to provide constructive comments and suggestions, and we believe that this will enable us to improve the quality of the analysis and to clarify the positioning of the study and the results that can be inferred from it.

We provide below our replies to each of the comments and outline our approach for revising the manuscript accordingly. Our replies are shown in blue, while the Reviewer's original text is reproduced in black. When we indicate line numbers, they refer to the revised manuscript with highlighted changes, unless otherwise specified.

Summary

This manuscript uses the Skywarden citizen-science database to study the occurrence characteristics of dune aurora, with the aim of constraining formation mechanisms. The scientific motivation is clear. The Skywarden community effort itself is exceptional: sustained and careful documentation by many aurora photographers has created a valuable record of rare and subtle auroral forms, and the authors' attempt to integrate this resource into research is timely.

My main concern is about what can be supported as “statistics” with this dataset. The manuscript acknowledges that biases exist, and I agree the authors are aware of that. However, in its current form the paper still presents “occurrence dependences” (MLT/season/activity/solar wind context) in a way that can easily be read as statistical properties of the phenomenon. Given the dataset structure, these quantities are not well-defined, and there is a risk that such results will later be cited as established occurrence trends even though they mainly reflect reporting and detectability effects. Below I outline the specific issues.

Thank you for this overall assessment and for pointing out the challenge of adequately presenting the results that can be obtained from the Skywarden observation analysis. We fully understand the concerns and have done our best to address them in the revised manuscript. In particular, one immediate measure to mitigate the risk of misinterpretation of our results in terms of occurrence of dunes rather than of dune reports has been to amend the title of the manuscript, removing “statistical” from it.

Major concerns

1) Occurrence inference without observational effort (a denominator)

To infer occurrence rates or systematic dependences, it is necessary to define the observational effort: when/where the phenomenon could have been detected, and under what observational

effort. Skywarden is excellent as a collection of reported cases, but in general it does not include the kinds of information that are easily available in scientist-run observations, such as:

- when observers were actually watching/shooting (active vs inactive time)
- how long and how continuously they monitored
- camera pointing, settings, and time resolution / cadence
- times when aurora was present but dunes were absent (true negatives), or present but not reported (missed/overlooked cases)

Without an effort model, or at least a defensible proxy, the results describe the distribution of reports rather than the occurrence properties of the phenomenon.

Thank you for raising this crucial point, which calls for expanding the discussion of what kinds of metadata are available from Skywarden. Some of the missing elements listed above can in fact be provided, at least to some extent, by Skywarden.

For instance, the active time of photographers is readily included in the dataset, as each report contains the information on when the observation started and ended. Within this time interval, we have estimated when the dunes could be reliably be deemed present (as described in the manuscript; see also the reply to Comment 2 below).

Concerning camera pointing, settings and image cadence, such elements can in many cases be retrieved. However, within the scope of our study, these do not necessarily bring insightful information: the dunes being a relatively long-lived phenomenon (compared to dynamical auroral forms such as spirals or rayed arcs, for instance), the temporal resolution of the imaging is not crucial to know. Similarly, since we do not aim at quantifying the brightness of the structures in this study, camera settings are not relevant either. We however acknowledge that the camera pointing direction could be used to better constrain the estimate of the dune location with respect to the observer. Nevertheless, our position for this study is to avoid undertaking complex image analysis and focus on the observation reports, leaving more elaborated mapping of the dunes for a future study (see also the reply to Comment 6 below).

We agree that it is presently not possible to reliably estimate the effect of missed dune cases on the results. In a way, this is a problem analogous to that, inherent to all the ground-based optical studies, of the effect of cloudiness when deriving occurrence rates of auroral phenomena from all-sky imager data spanning multiple dark seasons. We will never know how many events are missed due to cloudy conditions, and the fact that cloudiness has seasonal patterns certainly induces a bias in any statistical analysis of optical observations, which is challenging to quantify.

Concretely, what we can do within the scope of this study is make sure that what we report and discuss are presented as dune observation properties, that are not necessarily fully representative of the dune occurrence rates and properties. We have made sure to convey this message throughout the revised manuscript, adjusting the phrasing where appropriate (see the version with highlighted changes to easily identify where we have made modifications) and leave it for a future study to tackle a more in-depth analysis of dune images to get one step closer to assessing the intrinsic properties of dune aurora in itself.

2) Reporting and detectability biases are not bounded; normalization does not remove them

Dunes are low-contrast structures, and detectability depends strongly on observer behavior, equipment/settings, and retrospective inspection. Reporting probability is also likely to vary with local time, geomagnetic activity, and community attention. For example, once brighter and more dynamic aurora appear (e.g., breakup arcs), observers may change pointing, switch targets, or stop observing because they feel they have already captured what they wanted, which can imprint an apparent MLT preference even if the underlying phenomenon does not.

I understand the authors' approach of normalizing dune reports by the total number of aurora reports, with the goal of reducing sampling/reporting bias. However, dividing one biased distribution by another biased distribution does not, in general, remove bias. Many of the relevant biases (observer activity patterns, stopping rules, pointing choices, visibility/detectability, and social amplification) affect both the numerator and the denominator, and not necessarily in the same way. Even if some bias components are shared between the two, cancellation in a ratio only occurs under restrictive assumptions—essentially when the same multiplicative bias factor applies to both quantities in the same way and is independent of MLT, activity level, and other variables of interest. Without a way to quantify or bound these effects, the reported MLT/season/activity dependences cannot be interpreted as intrinsic properties of dune aurora.

We again agree with your arguments, although we wish to point out that we do not literally *divide* nor even normalise the dune event distributions by the distributions of all aurora observations. Rather, we simply compare the two and discuss reasons why they may differ.

That being said, we have improved the analysis of the MLT distributions of dune aurora and all-aurora reports by considering both start and end times of dune/all-aurora observations as inferred from the reports (instead of only the start times in the originally submitted version of the study). While this does not entirely remove the biases you point out, this new approach provides more robust results in the sense that the all-aurora MLT distribution are now more representative of the total observation time by Skywarden contributors. Interestingly, if we consider the MLT distributions (Fig. 3a,c), we note that they now both look closer to Gaussian distributions than with the prior analysis only considering the start times of observations (the distributions then looked less symmetrical). As is detailed in the revised manuscript, we have assessed the statistical significance of the differences between the two MLT distributions by carrying out permutation tests (see the newly added Sect. 3.4). The calculated difference of 2.1 h in MLT in terms of circular mean of those distributions is found to be significant, including after false discovery rate correction (see Table 1). A similar analysis has been done for the distributions associated with other considered parameters (monthly distribution of reports, driving parameters, and geomagnetic indices).

Nonetheless, given that it is not possible to properly quantify all the existing biases in the obtained distributions, our approach in the revision has also been to make it fully clear to the reader that we here report on dune *observations*, from which the information one can infer on

dune properties and occurrence is affected by biases at various levels. We have therefore expanded the discussion on those biases along the lines of your comment (l. 410–424; l. 540–542).

3) Event non-independence and potential double counting

Multiple submissions may refer to the same physical event, especially after community sharing triggers retrospective searches. The manuscript seems to be aware of this and attempts some handling, but there is no rigorous event-level consolidation. Without event clustering and verification (e.g., geometric checks using star fields), the effective sample size and the shapes of distributions remain uncertain. In that situation, distributions can be dominated by repeated sampling of a small number of events, which weakens any statistical interpretation.

We appreciate the comment and agree that the initial submission did not always guarantee that we avoid double counting of events. In the revised version, we have considered dune events rather than dune observations for the MLT analysis and the driving parameter analysis (Figs 3 and 5; see l. 188–193 for more details). We have however decided to keep the analysis at event level when it comes to the SME and SMR distributions (Fig. 6), given that those parameters can have fast variations.

Given that dunes are a fairly large-scale phenomenon (they can span hundreds of kilometres across), we think that our current criterion for considering observations to belong to two separate events is reasonable, although we identified a way to make it more robust: instead of looking at observation dates, we now look at the time interval between two observations, to cluster together observations occurring during the same night but on either side of UT midnight. In practice, this has led to absolutely no difference at all, but at least this is now handled properly from an algorithmic point of view.

4) Ambiguity in timing and duration definitions

Times are derived from mixed sources (EXIF and narrative descriptions). Remaining timing uncertainties are not quantified, nor is their impact on MLT assignment and on the OMNI/SME/SMR averaging windows assessed. Duration estimates are additionally affected by different cadence/continuity and variable “presence” criteria across observers; this is not a standard measurement-error setting. As a result, duration statistics provide limited physical constraint as currently presented.

While we agree that the accuracy and reliability of the time stamps obtained from mixed sources is not fully guaranteed, we already discuss with a reasonable amount of detail how we handled this (l. 93–98; l. 142–146; l. 195–199; l. 371–372; l. 380–l. 387 of the original manuscript). In particular, when the timing information could not be reliably confirmed, the corresponding observation was not retained for time-sensitive analysis (MLT, duration, associated driving conditions). Since in our analysis the data from separate observation reports are binned at no higher temporal resolution than one hour, we are confident that the vast majority (if not all) of the Skywarden reports retained in the analysis have sufficient time accuracy for our needs. Besides, when discussing the duration of observations and dune events, we stress that these are conservative low estimates, for the reasons that you highlight in your comment (l. 264–265).

5) Interpretation of solar-wind/geomagnetic parameter distributions

Comparisons between dune reports and all aurora reports are presented as evidence for preferred parameter ranges or structured distributions. However, given (i) unknown observational effort, (ii) reporting bias that varies with activity and time, and (iii) non-independence/double counting, fine-scale distribution features (e.g., dips near specific IMF values or apparent multimodality) are not compelling as physical selectivity. More importantly, this is a structural limitation of using Skywarden alone for “occurrence statistics”: without effort constraints and event independence, such distributions are hard to interpret physically and should not be presented as robust statistical properties.

After implementing the modifications in how we analyse the MLT distribution of observations, in how we avoid double counting of events, and in how we interpret the figures (i.e. emphasise that we are focusing on observation reports, which do not necessarily fully capture the intrinsic properties of dune aurora), we believe that this aspect has to a large extent been addressed. We have moreover significantly rewritten the interpretation of the solar wind and geomagnetic parameter distributions in the revised version, given that the permutation tests carried out on those distributions have revealed that the differences (often small in terms of retained statistics) are not statistically significant. This in itself is actually a very interesting result, which we have added to the conclusions (l. 633–637).

6) Latitude: observer location vs. dune location

The geomagnetic latitude used in the analysis is that of the observer rather than the dune emission region. If the goal is to understand dune behavior, the relevant quantity is the location of the phenomenon, not the location of the photographer. This conflation weakens latitude-related conclusions and also reduces the ability to de-duplicate events across sites. Since stars are visible in many images, approximate mapping/geolocation (with altitude assumptions) seems feasible in principle and could strengthen both dune-location inference and event consolidation.

We fully agree with this statement, and discussed it as a major source of uncertainty concerning the location of the dunes. We quantified the related uncertainty in terms of geomagnetic latitude (l. 398–405 in the original manuscript) and, following a suggestion made by Reviewer #2 (see the response to their Comment 11), we have similarly evaluated the location uncertainty in terms of longitude and MLT. Given that the phenomena we are discussing in this study (dune aurora, auroral electrojets) have fairly large characteristic scales (hundreds of kilometres in horizontal directions), the uncertainties in the real location of the dunes are not expected to significantly affect our results. We have added this point of discussion in the revised manuscript (l. 475–479). Besides, we have indicated with circles the possible locations of the dunes associated with individual observations used for the SECS analysis (Figs 7–9). We can see that, in all cases, despite the large (and on purpose conservative) uncertainty in the actual dune location, this does not affect our conclusions regarding whether the dunes have been observed in relation to the westward or eastward electrojet.

We are currently developing a method to map citizen science pictures of the aurora using the star field, but before being able to apply it to a large number of observations, it needs to be finalised, tested, and published. Within the scope of the current study, what we can do instead is estimate the impact of the uncertainties related to the discrepancy between dune observer location and dune aurora location and their effect on the calculated latitude and MLT distributions, as explained above. We are however looking forward to being able to carrying out the type of mapping that you suggest in future studies!

7) SECS/equivalent current section: selection criteria and leverage

The criteria for selecting events for the SECS analysis are phrased qualitatively (e.g., “several clear observations”, “a very large number of observations across a broad area”), which makes the selection hard to reproduce. I do think the point that dunes are not restricted to only westward electrojet context (i.e., they can also occur with an eastward electrojet) could be an interesting result. If the authors want to make that a key message, it may be stronger to present it through a smaller number of well-defined case studies with clear selection rules and deeper analysis, rather than embedding it within a broad statistical survey.

Thank you for this comment. We have consolidated the criteria for clustering individual observations together for the SECS analysis and modified its description in the text accordingly (l. 429–438). We have also improved the figures showing the three types of situations with respect to the eastward and westward electrojets by indicated the position of each observer (white stars) and adding circles quantifying the uncertainty in the location of the considered dune observations due to the dunes not being necessarily photographed at zenith. We have amended Figs 7–9, Table 2, and the text accordingly.

8) Discussion scope

The Discussion covers many mechanisms but would benefit from clearer prioritization tied to what the present dataset can and cannot constrain under these limitations.

Thank you for the suggestion. We have slightly revised the order in which we discuss the limitations and caveats of the Skywarden dataset. We have also added a few numbers to put the discussed points in perspective.

Minor comments

- Please define “solar wind density” explicitly (e.g., proton number density).
- Replace qualitative phrases such as “a very large number of observations” with quantitative values.
- Clarify how cloud okta is calculated in Appendix B and reconcile any inconsistencies with stated observation latitudes (observers latitudes vs. cloud latitudes).

Thank you for noting those points. We have implemented the changes as suggested in the revision. In particular, we now provide much more details on the cloud okta calculation in Ap-

pendix B, including a map of the stations making the cloudiness observation network (Fig. B1).

Overall assessment

Skywarden is extremely valuable for demonstrating the existence and diversity of dune aurora and for collecting candidate events. However, unless observational effort can be constrained and events can be consolidated at the event level, it is difficult to present the manuscript's "occurrence dependences" as statistical properties, and it is risky to leave them in the literature as a basis for physical interpretation. A clearer path would be either (i) to reframe the paper as a descriptive catalogue / case-based study and stop making claims about statistical occurrence characteristics, or (ii) to analyze a scientist-run dataset with well-defined observational effort if the goal is to find the statistical behavior of the phenomenon.

Thank you again for these excellent comments, which invited us to phrase more carefully the interpretation of our results. We have worked on the revision along the lines of what you suggest as option (i), in particular amending the title of the manuscript, improving the robustness of the analysis wherever possible, and elaborating on the limitations to fully clarify that our results cannot be interpreted as statistical properties of the dunes themselves. We hope that this revised version addresses your concerns and helps avoid the pitfalls you underlined in your comments.

Manuscript egusphere-2025-5374: Dune aurora: Statistical survey from a citizen science database

Response to Reviewer #2

We thank the Reviewer for their enthusiastic and thoughtful assessment of our manuscript and for raising important points to improve the paper. We truly appreciate the time and effort spent to provide constructive comments and suggestions, and we believe that this has enabled us to improve the robustness of the analysis and to clarify some points that were not made explicit enough in the original version.

We provide below our replies to each of the comments and outline how we have revised the manuscript accordingly. Our replies are shown in blue, while the Reviewer's original text is reproduced in black. When we indicate line numbers, they refer to the revised manuscript with highlighted changes.

Summary

This is an exciting and truly excellent piece of research. Dune aurora is a phenomenon which has only recently been reported in the scientific literature. This paper builds upon the earlier reports of this phenomenon by presenting the first statistical study of dune aurora. There are limited science-grade observations available at present, and the authors overcome this challenge by using citizen science observations of dune aurora. Using such observations is challenging and, barring a few points that I have highlighted within this review, the authors have met this challenge magnificently. The work is novel and exciting, and the paper is exceptionally well written. This work will be of great interest to the readers of *Annales Geophysicae*. On a personal note, I would like to add that I this is the most enjoyable paper which I have ever been asked to review.

Thank you very much for your very positive opening statement!

However, before I can recommend publication, there are some issues that need to be addressed.

Major comments

Comment 1: Lines 69-73: Are dune aurora more than just a projection effect?

Please would you expand upon the statement “the dunes have so far remained challenging to identify from all-sky camera images” to explain why it's so tricky for scientific instruments to capture an effect which amateur aurora hunters readily observe. More importantly, your comment that dune aurora “tend to be more visible when viewed at relatively low elevations” suggests that these are not a real physical phenomenon at all, instead they are simply a projection effect. I don't believe this to be the case (as there have been previous publications on dune aurora), but please would you clarify and revise this statement?

Thank you for this important comment. The reasons why the dunes are so difficult to identify

on all-sky images are likely multiple:

- if the emission is dim and the emission layer is thin (in altitude), it is probable that near-zenith the line-of-sight-integrated photons lead to a low signal, difficult to discern from the background;
- the dunes being relatively large structures (tens of kilometres between wave fronts) with moderately low contrast, the apparent diameter of a dune wave front overhead is significantly larger than if viewed at low elevation, which can make it difficult to identify luminosity gradients;
- in turn, when dunes are at low elevation (and hence should be easier to identify in optical data), the geometry of all-sky lenses tends to compress them close to the edges of the image, hindering a robust assessment of their presence.

We have elaborated on those points in the revision (l. 83–86), underlining that these are at this stage mere hypotheses.

Comment 2: Distribution of observations

Figure 2 shows that the observations were primarily made in the European sector, with a relatively small proportion coming from both North America and Australia / New Zealand. I appreciate and understand that a project which uses citizen science observations needs to use what is available, but some additional checks are needed to see how this distribution influences your results. What concerns me is that if there is a marked difference between these locations (due to, for example, changes in driving by the lower atmosphere) then it is possible that this could be responsible for some of the patterns observed (if, for example, all of the observations of dune aurora from the southern hemisphere were in October, then removing them may change both Figure 3 and the arguments which follow from it). Figures 3-7 need to be reproduced using just the observations from Europe. If the results remain unchanged, then this gives confidence that results are robust to the mix of locations. If the results do change, then this may shed some further light onto the driving processes.

I appreciate that you don't want to add another five figures to the paper, so I suggest that the additional versions of Figure 3-7 go into appendix C (unless they show something that fundamentally changes the results).

I do not suggest repeating the study with just the north American or Australia/New Zealand samples, as it seems likely that there are insufficient observations there to give robust results.

This is an excellent idea, which we gladly implemented in the revision. We have added the suggested figures in Appendix C and refer to them in the manuscript (l. 216–218; l. 276–277; l. 309–311; l. 356–373).

Comment 3: Lines 90-95: Consistency in the interpretation of images

You have explained that you visually inspected all of the events. I agree that this is the best way to approach a newly-reported phenomenon which is not fully understood, however steps need to be taken to ensure consistency in the interpretation of such images. For example, some authors

who have undertaken similar types of studies have done a second pass through all images to check that the same results are obtained (even if a small number of images are classified differently on the second pass). Other authors have used other approaches. Please would you explain what was done to ensure consistency and, if tests were not done, then these should be conducted.

Indeed, we agree that this point needed to be expanded upon, as it is a critical initial step forming the basis of the analysis. In fact, one may consider that the dune events were identified and validated at several levels by different people:

1. The observer themselves, in most cases, indicated that their pictures contained structures that may be dunes. Some experienced observers are familiar enough with the dunes to reliably identify them in their material, whereas others may be unsure about it but were encouraged to review their observations aided with freely available material providing guidelines for auroral form identifications (both integrated into Skywarden itself and as open material such as the ARCTICS Aurora Field Guide (Herlingshaw et al., 2024)).
2. The Skywarden moderating team review each report submission and, as part of this process, assess which auroral forms are present in the submitted material. During this step, the moderator updates the auroral form tags according to their knowledge and (extensive) experience in auroral form identification. This means that not only dunes identified in step 1 are approved or rejected by the moderator, but also that some observations where dunes were not identified by the observer were marked as containing dunes in this step.
3. For this study, two of the co-authors (MG and EB) have visually inspected all the Skywarden observation reports marked as containing dune aurora. In a few cases when either of them had reasonable doubt about the validity of the identification (e.g. suspicion of low-altitude clouds modulating diffuse aurora, suspicion of wave structures in airglow rather than dune aurora, barely discernible luminosity modulation in the images), the events in question were discarded.

We have added this information in the revised manuscript (l. 108–123).

Comment 4: Observer bias

You have correctly commented upon how human behaviour may influence your results (lines 156–157). However, there are two additional points that may be important:

- Amateur observers are more likely to go out to conduct observations when aurora is predicted, so it is likely that a greater number of people are observing when conditions are predicted to be disturbed.
- Once amateur observers have observed the aurora, and have been outside for a number of hours, they are likely to return their accommodation. Therefore, since dune aurora appear to be primarily associated with earlier MLTs than all aurora, dune aurora may be somewhat over-reported in the present study.

I believe that it is worth mentioning both of these points in lines 156-7, and the second of these points in the discussion.

Thank you for these suggestions. We fully agree that these are also important points to mention as possible biases from human behaviour. We have added them where suggested in the revised

manuscript (l. 203–206).

Minor comments

Comment 5: Lines 12-17: Level of detail

I think that these lines of the abstract could really do with some numerical details, so that the key findings are clear within the abstract. I suggest amending these lines as follows (providing all of the statements remain justified, given my comments under ‘major comments’). “We find that the vast majority (XX %) of dune observations take place within XXX degrees latitude of the equatorward boundary of the auroral oval. The majority (XX %) occur in the dusk sector ... ” etc.

Thank you for this comment; this is a good idea. We have added quantitative statements in the abstract of the revised manuscript.

Comment 6: Line 13: Clarity of text

I found the statement “in the dusk sector earlier than the peak in all auroral report distribution” to be unclear. This could be interpreted as ‘dune aurora occur in the dusk sector and earlier than the peak’ or ‘when dune aurora occur in the dusk sector (as opposed to the dawn sector), they are earlier than the peak’. A comma after “in the dusk sector” would resolve this ambiguity.

Point noted, we have added a comma in the revised text.

Comment 7: Line 155: Seasonal variation

You correctly state that there are ‘no events in June and July’. As this is winter in the southern hemisphere, it seems surprising that there were no reports of dune aurora from New Zealand or Australia. It would be worth commenting on this.

Thank you for noting this. Indeed, this is a point that should be discussed when interpreting the figure. The main reason behind the absence of reports from the southern hemisphere despite it being local winter conditions is that all but three southern-hemisphere dune reports to Skywarden come from the May 2024 superstorm (the remaining three are from the 11 October 2024 storm). The majority of the Skywarden reports from New Zealand and Australia were obtained thanks to the help of aurora chasing group community managers who posted our call for dune observations during the preparatory phase of this study. Since the May 2024 event was of exceptional magnitude, a great number of people photographed the aurora from those southern latitudes and discussed on social media about the dune display showing in many pictures. It is only very recently that Skywarden became known outside of the Nordic community, hence there are currently only a handful of reports from photographers based outside of Europe. We have summarised this in the discussion of the revised manuscript (l. 197–201).

We are extremely thankful to the aurora chasing group administrators for their strong engagement and support in getting Skywarden known in North America and Oceania, and we hope

that gradually the number of reports from these geographic areas will increase, possibly enabling more comprehensive studies of auroral processes in future.

Comment 8: Line 172: okta

Many readers will not be familiar with this unit (I had not encountered this previously). I appreciate that an explanation is given in the first three lines of appendix B, but I think an explanation is also needed in the main part of the manuscript.

Thank you for the suggestion, we have implemented it in the revision (l. 223–226). We have also, following a comment by Reviewer #1, provided more details on the cloud okta calculation in Appendix B.

Comment 9: Lines 230-260: Results & statistics

This section of the manuscript draws comparisons between the ‘dune observations’ and ‘all aurora observations’. I believe that this analysis would be strengthened by conducting statistical tests to show that the distribution of the ‘dune observations’ is not consistent with these being drawn from the ‘all aurora observations’ sample in each case. This a suggestion, not a requirement, as the counter-argument is that detailed statistical tests may overly complicate this section of the paper.

Thank you so much for suggesting this excellent idea! We decided to carry out statistical tests to provide more robust grounds for discussing the discrepancy in peak MLT for dunes vs all-aurora observations (as well as other parameter distributions). The suitable approach in our situation (parameter distributions associated with 61 dune events, to be compared to nearly 10,000 of aurora reports in Skywarden) is to carry out permutation tests, for which we consider a statistical metric (e.g. circular mean for MLT, median value for the driving parameters, total variation distance for the monthly distributions) and calculate the p -value of the effect, with 10,000 iterations of a permutation test. Very interestingly, the results (presented in Table 1) indicate that the differences in dune-event vs all-aurora distributions presented in Figs 3, 5, and 6 are statistically significant (including after false discovery rate correction) for MLT, monthly occurrence, SME, and SMR. On the other hand, the differences between distributions for the driving (solar wind and IMF) parameters are *not* statistically significant. We describe the procedure and discuss the obtained results on l. 383–409.

We believe that these are very interesting results that greatly enhance the impact of our study. Thank you again for this great suggestion!

Comment 10: Line 296: Interpretation of Figure 7

The text states that ‘both curves [the magenta and black lines in Figure 7b] are relatively close to each other’. I would disagree with this statement, as the magnitude of the gradient varies by a factor of 2. I also note that, if the data points from the southern hemisphere were not considered, the gradient of the black dotted line would change, bringing it closer to Weyland et

al. (2023). I suggest that the interpretation of this plot is re-visited once comment 2 has been addressed.

Thank you for pointing this out; we agree with your argument. In fact, considering that the figure in question was bringing little to the overall study, we have decided to remove it (together with the related discussion) in the revised manuscript.

Comment 11: Lines 398-405: Longitude variations

You correctly comment that the latitude of the dune aurora is within ± 4.2 degrees of the observer. However, a similar comment is needed for the uncertainty in longitude & MLT. This is not needed for every observation (calculating it for every latitude would be somewhat tedious), but simply giving a couple of ballpark values would be sufficient. Perhaps giving a range of longitudes and MLTs for an observer at 70 GLAT and at 60 GLAT would be enough. Of course, if this range of values causes uncertainties about whether dune aurora are in the eastward or westward electrojet (Table 1), or if the uncertainty in the MLT of the observations is larger than the changes in MLT between the ‘dune observations’ and the ‘all auroral observations’ (Figure 3), then this also needs to be discussed.

Thank you for this comment and related suggestions. We have added an estimate of uncertainties in terms of longitude (l. 535–537). Moreover, in the SECS analysis figures, we have added markers (white stars) indicating the positions of individual observers, around which we have drawn circles rendering the uncertainty in the dunes’ real location (radius of 463 km, consistent with the value given on l. 474).

We note that these longitude uncertainties ($\pm 8.3^\circ$) correspond to approximately ± 0.55 h in MLT. We have added this point to the discussion too.

Additional comment

Lines 439 – 445 discuss gravity wave activity, and the filtering of such waves. Such waves can be launched by orography, thunderstorms, fronts, baroclinic instabilities in the troposphere, volcanic eruptions, and the excitation of waves associated with the polar vortex (Fritts and Alexander, 2003). The Northern Annular Mode (NAM) index is a measure of the strength of the polar vortex (Kumar et al., 2025; Gerber and Martineau, 2018), and I wonder if a statistical link could be found between this index and the occurrence of dune aurora or whether a case study could be conducted to look for a link between dune aurora and gravity waves launched by the polar vortex. This is probably beyond the scope of the present study but, if you are interested in pursuing this idea, I would recommend contacting Corwin Wright, Neil Hindley and Juliana Jaen at the University of Bath who, in my opinion, would be exceptionally well-placed to help with this work.

Thank you for these very concrete suggestions, we appreciate them and take good note of them in view of the follow-up dune studies. We have included some of those points into the discussion (l. 582–589), emphasising that the possible role of atmospheric dynamics is, at this stage, a mere

speculation.

References

- Fritts, DC, Alexander, MJ. 2003: Gravity wave dynamics and effects in the middle atmosphere. *Rev. Geophys.*, 41: 1003, <https://doi.org/10.1029/2001RG000106>
- Gerber, EP and Martineau, P. 2018. Quantifying the variability of the annular modes: reanalysis uncertainty vs. sampling uncertainty, *Atmos. Chem. Phys.*, 18: 17099–17117, <https://doi.org/10.5194/acp-18-17099-2018>
- Kumar, A, Stolle, C, Yamazaki, Y, Pedatella, NM, Kunze, M, et al. 2025. Impact of weak and strong stratospheric polar vortices in the northern and southern hemispheres on solar-migrating semidiurnal tides in UA-ICON. *J. Geophys. Res. Atmos.*, 130: e2025JD043550. <https://doi.org/10.1029/2025JD043550>

Summary

This is an excellent piece of work which will make a valuable contribution to the research field. Thank you very much for your positive assessment and constructive comments and suggestions!