

We thank the referee for the careful and constructive review and for the positive assessment of the manuscript. The comments helped us clarify the interpretation of IM, improve the presentation of Figure 2, examine possible observational biases in the site-to-site comparison, and expand the discussion of microbursts.

All line numbers cited below refer to the revised manuscript with tracked changes, rather than to the clean manuscript.

Comment 1

Line 58: The text states that the IM “presumably reflect the generation of discrete chorus elements”, implying that patchy pulsations do not contain such discrete elements, but I think this is a mischaracterization and is not supported by the data included in this study. The patchy pulsating aurora (without IM) may well contain discrete elements but only lower energy precipitation (possibly from the lack of high-latitude propagation, as argued in Miyoshi 2020). In that 2020 paper, the ‘patch’ luminosity varies on the larger scale for lower energy precipitation while the bursty luminosity, that corresponds to individual chorus elements, varies on the scale of IM. It doesn’t mean there is no generation of chorus elements just because there is no IM present. I think this can be more plainly stated here in the introduction, by making it clear that there are multiple reasons why IM would or would not be present — as are discussed in the Discussion section.

Response:

We agree. We have revised the final paragraph of the Introduction so that the absence of optically identifiable IM is not interpreted as evidence for the absence of discrete chorus elements. We now state that multiple factors may affect whether IM is visible in optical observations, including variations in the temporal structure and propagation of chorus waves. This replaces the earlier statement that the conditions favoring IM presumably reflected the generation of discrete chorus elements (Lines 58–61).

Comment 2

Figure 2 and associated text: It would be excellent to see a line plot of the IM, as well as the larger patches. The IM are difficult to pick out in the keogram format, but a horizontal cut would show the modulations clearly as well as how close to background they are, which is important for reasons stated later in the paper. I suggest adding a panel showing these line plots for panels (a) and (b).

Response:

We agree. We added panels (a’) and (b’) below the corresponding keograms in Figure 2. Each line plot shows the raw counts averaged over the narrow horizontal band indicated by dashed lines in the keogram. The two line plots use the same vertical span in counts, allowing the modulation amplitudes to be compared directly. Panel (a’) shows the relatively gradual intensity variations associated with the larger-scale patches, whereas panel (b’) makes the rapid IM and its intensity relative to the background much clearer. The associated description was revised accordingly (Lines 136–151 and the revised Figure 2 caption).




Figure R1. Revised Figure 2, including the line plots in panels (a') and (b').

Comment 3

Figures 4, 5, and Results: It is difficult to understand the physical reasons underlying the drastic difference in occurrence rate of IM over a 1.7 degree change in magnetic latitude. Could you speculate on this more in the Discussion? Is there no chance it is due to observational bias? You have done a rigorous job of explaining the robustness of the model results, but it is still a bit baffling that such a strong delta exists over this relatively small latitudinal range.

Response:

We agree that the site-to-site contrast requires careful evaluation and should not be interpreted as an unambiguous causal gradient with magnetic latitude. We therefore performed additional analyses and revised the Machine learning detections, Results, Discussion, Conclusions, and an added Appendix B. First, to reduce possible site dependence in the color-camera pre-screen, we broadened the original Diffuse-only criterion to $p(\text{Arc}) + p(\text{Discrete}) + p(\text{Diffuse}) \geq 0.8$ and clarified that the color-camera output is used only as an aurora-or-not pre-screen. This modification was intended primarily to reduce the effect of differences in camera type and sensitivity settings, which could make the boundary between the Discrete and Diffuse classes site dependent. PsA and IM are subsequently determined from the qCMOS keograms (Lines 104–124). The revised analysis still shows positive Kiruna–Skibotn differences of 0.293, 0.390, and 0.221 at 4, 5, and 6 MLT, respectively. The strict sensitivity ranges exclude zero at 4 and 5 MLT and marginally include zero at 6 MLT; the weighted sensitivity ranges exclude zero throughout 4–6 MLT (Lines 291–296 and 366–370).

We also examined the availability of raw qCMOS data, observable exposure, differences in the color-camera pre-screen, lunar conditions, weather, and the occurrence of PsA itself (Appendix B, Lines 524–565). For bins with at least 5 h of observable exposure, PsA occupied 42.8% and 24.8% of the observable time at Kiruna and Skibotn, respectively, while IM occupied 37.3% and 26.6% of the PsA-positive time. Thus, unequal observational opportunity explains part of the difference in the total number of detected events, but a difference in the conditional IM fraction remains after normalization by both observable exposure and PsA occurrence. The qCMOS cameras at the two sites had the same optical and acquisition setup. Residual effects of weather and the color-camera pre-screen cannot be completely excluded.

As a possible geometrical and physical explanation, we now discuss that the approximately 1.6° MLAT separation may cause the two sites to sample systematically different parts of the auroral oval. Grono and Donovan (2020) found a 1–2° MLAT separation between the occurrence peaks of different PsA morphologies, demonstrating that auroral properties can vary over a comparable latitudinal scale. We therefore interpret the result as a robust site-to-site difference associated with MLAT, particularly at 4–5 MLT, rather than proof of an intrinsic causal MLAT dependence (Lines 371–383 and 493–498).




Figure R2. New Figure B1, showing observable exposure, the fraction of observable time classified as PsA, and the fraction of PsA time classified as IM at Kiruna and Skibotn.

Comment 4

Discussion: I think it might be worth commenting even more on microbursts directly and their potential relationship to the conclusions presented here. If IM are indeed a signature of the high-energy tail of PA, then the kind of statistics you show here is very important for understanding radiation belt dynamics and that energy transfer from the magnetosphere to the atmosphere. It would be great to include the observations by Shumko et al. 2021 and by Elliott et al. 2022 when discussing these implications. I think it's worth pointing out how ELFIN or REAL cubesats may help make the connection between the optical signatures you show and the microburst/relativistic signatures measured in the precipitation.

Response:

We agree and have expanded the Discussion. We now describe the strong correlation reported by Shumko et al. (2021) between patchy aurora and >1 MeV electron microbursts, while noting that the optical cadence in that study was insufficient to determine whether IM was present (Lines 448–451). We also discuss the multipoint study by Elliott et al. (2022), which constrained the spatial extent and duration of a microburst-producing chorus region and demonstrated the possible importance of microbursts for outer-radiation-belt electron loss (Lines 455–460). Finally, we point out that coordinated optical observations and low-Earth-orbit particle measurements, including ELFIN observations and measurements from the recently launched REAL CubeSat, could test the connection between optical IM and the energy, pitch-angle distribution, and temporal structure of microburst precipitation (Lines 463–465). We retain the qualification that the optical observations alone cannot identify the IM-dominant cases as relativistic microbursts.

We thank the referee again for the constructive comments, which substantially improved the clarity and scope of the manuscript.

We thank the referee for the positive assessment of the manuscript and for the helpful comments. We have revised the manuscript and carried out additional checks as described below.

All line numbers cited below refer to the revised manuscript with tracked changes, rather than to the clean manuscript.

Minor comment 1

l. 41: Would it be possible to indicate an order of magnitude for the typical amplitude of the internal modulation?

Response:

Yes. We added an order-of-magnitude reference in the Introduction. Røyrvik and Davis (1977) reported that the amplitude of the 3 ± 1 Hz modulation could reach approximately 20% in relative intensity (Lines 41–42). We describe this as a reported upper amplitude rather than a universal typical value, because the measured amplitude depends on the background intensity and observational conditions.

Minor comment 2

l. 96: The definition of the F1 score could be introduced already here (rather than on l. 131)

Response:

We agree. The definition of F1 has been moved to its first appearance in Section 3.1 (Lines 98–99). The later duplicate definition in Section 3.2 has been deleted (Lines 160–161).

Minor comment 3

l. 111: Would it be worth indicating how many different nights of observations comprise keograms labelled as containing PsA and as containing IM within PsA? It could be insightful to relate these numbers to that of manually annotated Kiruna keograms given just above.

Response:

We agree. We now report that the manually annotated dataset comprises 12,429 keograms from 98 observing nights. Of the non-ambiguous samples, 8,163 PsA-positive keograms were obtained from 90 nights, and 2,608 IM-positive keograms were obtained from 72 nights (Lines 131–134).

Minor comment 4

Fig. 2: Since this is the first figure showing keograms and discussing the PsA morphologies, I suggest adding arrows in the various panels to help the reader recognise the identified signatures, and refer to these in the text when describing the figure.

Response:

Thank you for the suggestion. We added arrows to panels (a)–(c) to identify representative PsA structures without IM, regions in which IM is visible, and a discrete auroral structure that is not classified as PsA. We refer to these arrows explicitly in the revised text (Lines 144–151). In addition, following the suggestion of Reviewer #1, we added line plots in panels (a') and (b'), which further clarify the distinction between the larger-scale patches and rapid IM (Lines 136–149 and the revised Figure 2 caption).




Figure R1. Revised Figure 2, including the added arrows and line plots.

Minor comment 5

l. 166: I would suggest mentioning the number of test keograms earlier in the manuscript (e.g. in Sect. 3.1) and briefly explaining how the test set was selected (in a similar way as for the training set) and how many different nights it contains.

Response:

We agree. We now define an observing night as 14:00 UT to 09:00 UT on the following day and explain that the 98 nights were randomly divided at the night level, so that no night was shared among the training, validation, and test subsets. The fixed split contains 9,398 keograms from 78

nights for training, 983 keograms from 10 nights for validation, and 1,805 keograms from 10 nights for independent testing. The test set contains 1,222 PsA-positive keograms from nine nights, including 432 IM-positive keograms from nine nights. We also state explicitly that the test set was used only for final performance evaluation (Lines 165–170).

Minor comment 6

I. 258: Was there any particular reason to select this specific night (20–21 November 2025)? Is there a way to assess to what extent it is a representative night? If possible, please briefly comment on it.

Response:

We selected this night because auroral activity persisted throughout the night, allowing the evening-to-morning evolution to be illustrated continuously, and because the observations from this night were not used for model training. We have added this selection criterion to the revised manuscript (Lines 305–306). We do not claim that this night is statistically representative; instead, Figures 7 and 8 are presented as an event-level illustration of the range of morphologies underlying the statistical tendencies.

Minor comment 7

I. 387: On microbursts, Kang et al. (2022) studied scale size and effects of oblique propagation of chorus waves in the inner magnetosphere; their results might prove worth briefly discussing in this paragraph.

Response:

We thank the referee for drawing our attention to this study. We considered adding Kang et al. (2022), but decided not to include it in the revised manuscript. That study addresses the scale size of microburst-producing chorus regions and the effects of oblique chorus propagation, whereas the paragraph in question follows a more specific argument linking the optical visibility of IM to the high-energy tail of PsA precipitation and to direct aurora–microburst observations. We found that inserting the citation without a separate discussion of its scale-size analysis interrupted this argument, while a proper treatment would require broadening the scope beyond what can be constrained by our optical data. We have therefore left this paragraph unchanged with respect to Kang et al. (2022).

Minor comment 8

I. 397: Could the claims based on Fig. 8B, regarding the morphology of the IM in the midnight and morning sectors, be substantiated with some numbers? For instance, you could make an argument based on the test data set that you have inspected thoroughly and which is of a reasonable size to make rough statistics, which would presumably not require an extensive amount of additional work.

Response:

We agree and quantified this tendency using the visually reviewed, IM-positive test keograms. After excluding cases with unclear morphology, including ambiguous cases that exhibited both a clear main pulsation with embedded IM and IM-dominant structuring, 410 keograms remained. The fraction for which the main pulsation was indistinct while IM remained clear was 36% at 19–01 MLT and 69% at 04–06 MLT. We added these values to the Discussion to substantiate the qualitative comparison between Figures 8B and 8D (Lines 472–474).

Minor comment 9

Discussion: I wonder whether from your analysis you could identify any seasonal trend for IM occurrence; is this something you have considered looking into? It could be worth briefly discussing it if there is anything worth noting on that respect.

Response:

We examined the monthly IM fraction within PsA using the same classification scheme and uncertainty estimates as in Figure 5; the result is shown below for the referee. No clear or consistent seasonal dependence is evident at either site. Because the dataset covers only two to three observing seasons and the monthly coverage is uneven between the sites, an apparent monthly variation could reflect weather or sampling differences rather than a reproducible seasonal effect. We therefore chose not to add a seasonal interpretation to the manuscript, in order to avoid over-interpreting the present dataset.




Figure R2. Monthly IM fraction within PsA for (a) Kiruna and (b) Skibotn. The classification and uncertainty estimates follow the same approach as Figure 5.

Minor comment 10

Discussion: It could be worth discussing whether IM may be preferably associated with patchy pulsating aurora (PPA) or amorphous pulsating aurora (APA), if this is an aspect that you could investigate without too many difficulties. Your results could then be discussed in relation to PPA/APA studies (e.g. Grono & Donovan, 2020; Tesema et al., 2020).

Response:

We agree that the relationship between IM and PPA/APA is an interesting topic. However, the present classifier was designed specifically to identify PsA and the presence or absence of IM; it was not trained to classify PsA systematically into PPA and APA. Moreover, the one-dimensional keograms used as model input do not always retain enough two-dimensional morphological information for a reliable PPA/APA classification. We therefore consider that adding a preference for either morphology without a dedicated classification analysis would be speculative, and we have not added such a claim to the manuscript. Grono and Donovan (2020) is now cited in a separate discussion of the latitudinal distributions of PsA morphologies (Lines 374–379), but not as evidence for an IM–PPA/APA association. A dedicated morphological classification will be a subject for future work.

We thank the referee again for the helpful comments, which improved both the methodological description and the interpretation of the results.