Reviewer comments 1 and 5:

The main conclusion, namely the MLT-dependent hardening of diffuse auroral precipitation, has already been reported in previous optical, radar, and satellite studies. The novelty therefore appears mainly methodological rather than scientific. In addition, the Bayesian regression quantifies a trend that is already visible in the data, without providing substantial additional physical insight.

Response:

We thank the reviewer for these important comments, which we address together. We agree that the broad physical picture of harder diffuse/pulsating auroral precipitation toward the morning sector is not entirely new, and we do not intend to claim that the present study reveals a fundamentally new precipitation mechanism. We will revise the manuscript accordingly to avoid overstating the physical novelty.

At the same time, we do not think that the contribution of the present work is merely methodological. The novelty of the manuscript is to provide a quantitative and statistically robust *spectroscopic* confirmation of a previously suggested tendency using a large, homogeneous ground-based dataset spanning two winter seasons. In this sense, the present paper is not intended as a technical note, but as an observational study that places stronger statistical constraints on a known trend.

Previous studies already provided important evidence for morning-side hardening, but some limitations remained. Hosokawa and Ogawa (2015) used 21 nights of radar observations. Kawamura et al. (2020) analyzed 37 nights of PsA observations and showed a statistically meaningful lowering of the emission altitude toward the morning sector, especially after 06 MLT, but without an explicit statistical determination of a preferred breakpoint range. Partamies et al. (2017) also reported MLT-related differences in pulsating aurora occurrence and altitude behaviour, but the transition timing remained less clearly constrained. This also clarifies the role of the Bayesian framework. We agree that the hardening is already visible qualitatively in the binned results, and the Bayesian analysis is not meant to "discover" an otherwise invisible trend. Its role is rather to quantify the robustness of that trend in a broad, non-Gaussian dataset with large dispersion and mixed diffuse populations, where a simple linear regression would not be sufficient. More specifically, it allows us to compare competing piecewise descriptions of the MLT dependence in a consistent way.

Following the second referee's suggestion, we developed a more robust Bayesian framework, which complements the analysis presented in the submitted manuscript. The goal of this additional analysis is to test where the transition toward higher inferred electron energy is best represented, while accounting for night-to-night variability. We will add a hierarchical Bayesian breakpoint analysis in which each data point is the median inferred energy proxy computed for one observing night and one MLT bin, as shown in Figure 1. This analysis is supported by the coverage of the central MLT range: each MLT bin from 22 to 6 MLT contains more than 20 independent nights, with typically more than 1000 spectra per bin. The coverage is strongest between 0 and 5 MLT, where each MLT bin contains 34 to 38 nights and about 1900 to 2600 spectra. The edge bins, especially before 21 MLT and after 7 MLT, have lower coverage and are therefore interpreted more cautiously.

The model first defines the expected energy proxy for a given night n and MLT bin m as

$$E_{n,m}^{\mathrm{model}} = \beta_0 + a_n + \beta_1 x_m + \beta_2 \max\left(0, x_m - x_{\mathrm{bp}}\right), \tag{1}$$

where x_m is the MLT coordinate and x_{bp} is the candidate breakpoint. Before the breakpoint, the MLT slope is β_1. After the breakpoint, the slope becomes $\beta_1 + \beta_2$. The parameter β_2 therefore measures the change in slope after the candidate breakpoint. The term a_n is a night-specific offset. It accounts for the fact that some nights are globally more energetic than others, independently of the common MLT trend. These offsets are estimated hierarchically as

$$a_n \sim \mathcal{N}\left(0, \sigma_{\mathrm{night}}\right). \tag{2}$$

The observed quantity fitted by the model is the median energy proxy $E_{n,m}^{\mathrm{obs}}$, computed for night n and MLT bin m. It is assumed to scatter around the model prediction as

$$E_{n,m}^{\mathrm{obs}} \sim \mathcal{N}\left(E_{n,m}^{\mathrm{model}}, \sigma_{\mathrm{res}}\right). \tag{3}$$

This model differs from the initial global MLT-bin analysis presented in the submitted manuscript. In the original analysis, all nights were collapsed into one median value and one MAD-based dispersion for each MLT bin. In the hierarchical model,

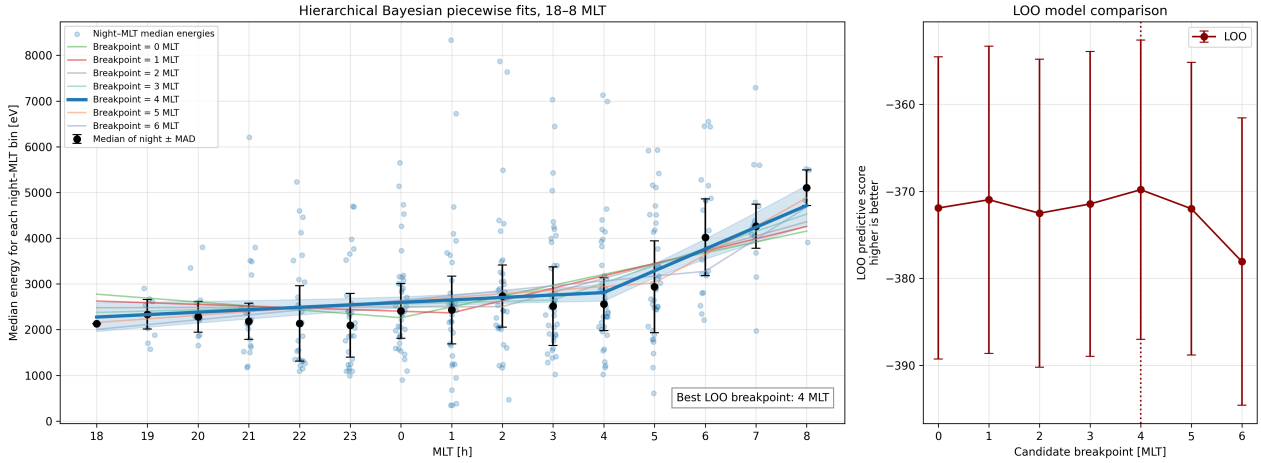

Figure 1. Hierarchical Bayesian breakpoint analysis using night–MLT median energies. Left: piecewise fits for several candidate breakpoints. Right: model comparison using leave-one-out cross-validation (LOO), showing the strongest predictive support for a breakpoint near 4 MLT, although neighbouring values remain statistically plausible.

the fit is instead performed on the median energy proxy computed separately for each night and each MLT bin. The model therefore keeps the night-by-night structure of the dataset, rather than reducing it to a single global MLT profile

This formulation allows us to separate two levels of variability. The first one is the night-to-night variability, described by σ_{night}. It measures how much the overall energy level can shift from one observing night to another through the night-specific offsets a_n. The second one is the residual variability, described by σ_{res}. It measures the scatter of the median value inferred energy proxy computed for one observing night and one MLT bin around the fitted piecewise MLT trend, after the night-to-night offsets have been taken into account.

The model was then fitted for several fixed candidate breakpoints, from 0 to 6 MLT. For all tested breakpoints, the inferred night-to-night baseline dispersion is substantial, with $\sigma_{\mathrm{night}} \simeq 1.0$ keV. This shows that systematic differences between observing nights are large and must be accounted for. At the same time, a residual scatter of about $\sigma_{\mathrm{res}} \simeq 0.85$ keV remains around the fitted MLT trend. This means that significant variability remains even after correcting for the night-to-night offsets and for the common MLT dependence. This is expected for a diffuse auroral dataset that mixes different activity levels and optical subtypes. The hierarchical analysis therefore does not simply reproduce a visually apparent trend. It also quantifies the level of variability within which the MLT-dependent hardening is detected.

We then compared the candidate breakpoint models using leave-one-out cross-validation (LOO). In this approach, each model is tested on its ability to predict individual data points that were left out of the fit. Here, these data points are the median energy proxies computed for one observing night and one MLT bin. The highest LOO score is obtained for a breakpoint near 4 MLT. However, neighbouring breakpoints remain statistically plausible. We therefore do not present 4 MLT as a sharp threshold, but rather as the preferred representative breakpoint within the tested range.

In summary, we agree that the manuscript should not be presented as introducing a new physical mechanism. However, we do think that it brings scientific novelty in the form of (i) a large homogeneous spectroscopic dataset, (ii) a statistically robust confirmation of the morning-sector hardening, and (iii) a first quantitative determination of the preferred breakpoint range within a Bayesian model-comparison framework.

Reviewer comment 2:

The manuscript groups together multiple types of diffuse aurora (...) into a single category based on AI classification. (...) I strongly recommend that the authors perform a refined classification within the "diffuse" category (e.g., isolating pulsating aurora), and re-evaluate the MLT dependence of electron energy separately for each subtype.

Response:

We thank the reviewer for this important comment. We agree that the diffuse aurora class used in the present study includes multiple optical morphologies and therefore potentially multiple precipitation regimes. We also agree that this limits the uniqueness of the physical interpretation, and we will revise the manuscript accordingly to make this point more explicit and to adopt a more cautious discussion of the possible mechanisms. At the same time, we note that the present study was intentionally designed as a statistical investigation of *diffuse aurora*, as already reflected in the title and scope of the manuscript, rather than as a subtype-pure study of pulsating aurora alone.

The main practical limitation is that the AI-based auroral classification used here is derived from all-sky images acquired at a cadence of one image per minute with an 8 s exposure time. This is sufficient to separate broad classes such as arc, discrete, diffuse, or no aurora, but not to robustly distinguish, within the diffuse class, clearly pulsating, patchy, amorphous, or related subtypes over the full two-winter interval. Therefore, even if a subtype-level separation were physically meaningful, the presently available classification product does not contain the temporal information required to perform such a re-analysis in a homogeneous way.

More importantly, we think that the physical mapping between optical morphology and wave mode is more complex than a strict subtype-to-mechanism correspondence would imply. As shown by Grono and Donovan (2018), the morphology of pulsating/diffuse aurora itself evolves with magnetic local time: amorphous pulsating aurora (APA) tends to dominate earlier, whereas patchy pulsating aurora (PPA) and patchy aurora (PA) become more common later in MLT. Therefore, even within pulsating aurora, morphology is not fixed across local time, and it is difficult to treat a single optical morphology as the unique signature of a single precipitation regime.

In addition, recent Arase-based studies indicate that the relation between optical morphology and wave mode is not one-to-one. Fukizawa et al. (2018) reported that not only lower-band chorus but also ECH wave intensity can correlate with pulsating auroral emission, explicitly identifying ECH as a candidate contributor to PsA. Fukizawa et al. (2020) further showed that ECH waves can scatter electrons of about 5 keV into the loss cone and contribute to weak visible auroral emission, with a simulated green-line intensity of about 200 R. We therefore agree that the reviewer's interpretation may be plausible as a first-order picture, but we feel that the stronger implication of a strict correspondence between morphology and scattering wave mode is not established in the recent literature. To our knowledge, there is also no recent ECH-based explanation for a systematic dawnside hardening comparable to that discussed for chorus-related precipitation. We therefore agree that part of the observed MLT-dependent hardening may reflect changes in the relative occurrence of diffuse auroral morphologies with MLT, rather than the evolution of a single uniquely identifiable process. The revised manuscript will be softened accordingly.

We note that this heterogeneity is also reflected in the hierarchical Bayesian analysis described in response to comments 1 and 5. The inferred night-to-night baseline dispersion is of order (1 keV), and a residual scatter of order (0.85 keV) remains around the common MLT trend. We therefore do not interpret the hardening as the behaviour of a single, subtype-pure precipitation regime.

Following the reviewer's recommendation, we nevertheless explored whether clearly pulsating intervals could be isolated using an independent high-cadence optical dataset. This was precisely the motivation for examining the EMBLA narrow-field camera (Whiter et al., 2024). EMBLA keograms provide a horizontal sampling of about 2 s per pixel, allowing positive identification of intervals with marked pulsating behaviour. We therefore developed a dedicated manual selection procedure from EMBLA keograms to identify intervals with clear pulsating signatures, as illustrated in Figure 2.

This EMBLA-based analysis remains exploratory and cannot support a homogeneous subtype-resolved statistical reanalysis of the full ASIS dataset. EMBLA is only available from mid-January 2025 onward, whereas the ASIS dataset analyzed here spans October 2023 to March 2024 and October 2024 to March 2025. The overlap is therefore limited to only about ten nights. In addition, non-detection of pulsations cannot be interpreted as absence of pulsating aurora, especially for faint events. Finally, some of the few available cases in key MLT sectors, especially around 23 MLT and 02 MLT, correspond to high-energy outliers and are not representative of the full diffuse-aurora population within those hourly MLT bins, as displayed by Figure 5 from the submitted paper.

Nevertheless, the EMBLA-selected PsA-only subset provides a useful consistency check. On the nights in common, it shows a qualitatively similar increase of inferred energy toward the morning sector as the corresponding all-diffuse subset, while the full diffuse dataset preserves the same global behaviour (Figure 3). We do not interpret this limited sample as a statistically

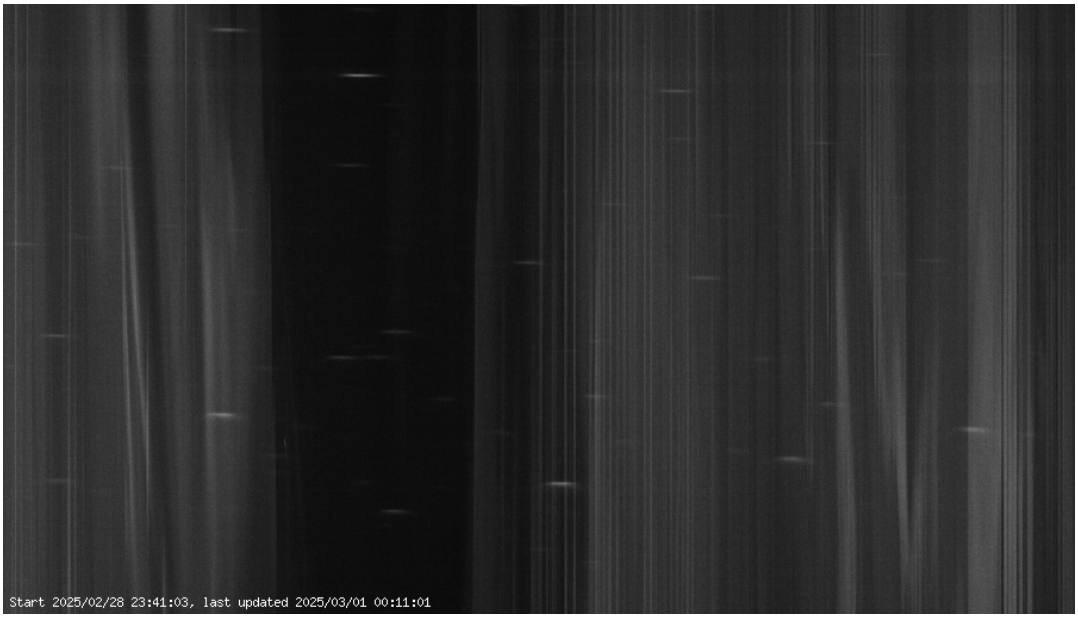

Figure 2. Example EMBLA keogram for 28 February 2025 showing a clear interval of pulsating aurora. The horizontal sampling is approximately 2 s per pixel. Courtesy of D. Whiter.

complete subtype-resolved analysis, but it suggests that the main dawnside hardening found in the diffuse dataset is not solely produced by subtype mixing within the diffuse class.

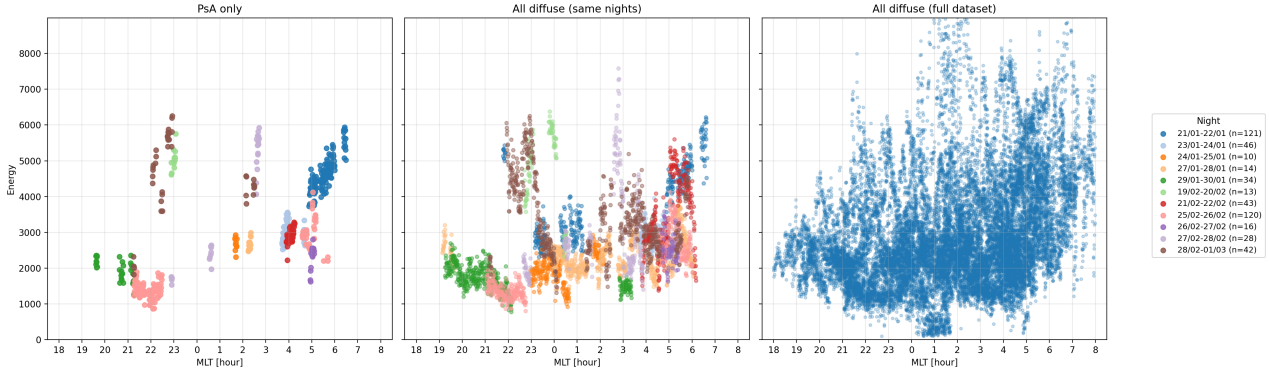

Figure 3. Comparison of inferred energy as a function of MLT for the EMBLA-selected PsA-only subset (left), the full diffuse sample restricted to the same nights (middle), and the full diffuse dataset (right). The PsA-only subset remains limited and non-homogeneous, but shows a qualitatively similar dawnside hardening tendency.

For these reasons, we do not think that a subtype-resolved statistical reanalysis of the full two-winter dataset would be robust at this stage. Instead, we will revise the manuscript to clarify that the present results apply to the diffuse-aurora population selected by the classifier, with pulsating aurora forming an important but non-exclusive subgroup. We will also soften the interpretation accordingly, and discuss the observed hardening more cautiously as a population-level property rather than as the signature of a uniquely identified mechanism. If the reviewer and editor consider it useful, we can also add a short discussion

paragraph in the manuscript on this point; however, we would prefer not to include the EMBLA-based analysis in the main paper itself, since it remains exploratory and not statistically representative of the full ASIS dataset.

Reviewer comment 3:

The inferred "mean energy" relies on a red/blue line ratio, a Maxwellian assumption, and steady-state transport models. However, diffuse and especially pulsating aurora are known to involve non-Maxwellian distributions, time-dependent precipitation, and multi-component electron populations. This raises a fundamental question: what physical quantity does the derived "mean energy" actually represent? The manuscript should more clearly discuss systematic biases in the energy estimation, possible MLT-dependent biases, and limitations in interpreting absolute energy values.

Response:

We thank the reviewer for this important comment. We agree that the inferred quantity is model-dependent and should not be interpreted as the true mean energy of the precipitating distribution in a strict kinetic sense. However, this limitation is inherent to optical line-ratio methods and not specific to the present study. In practice, such approaches are widely used in auroral studies together with transport models, including GLOW/ModGLOW-, Transsolo-, or TREx-ATM-type frameworks, all of which require an assumed incident distribution and are therefore intrinsically model-dependent (Adachi et al., 2017; Robert et al., 2023; Liang et al., 2024).

A full quantification of systematic biases, including possible MLT-dependent biases, would require a dedicated validation study combining forward modelling and independent constraints such as satellite particle measurements, which is beyond the scope of the present statistical analysis. In the revised manuscript, we will therefore clarify that the retrieved quantity is better understood as an *effective characteristic energy* (or energy proxy) inferred from the observed red/blue ratio within a specified transport-model framework, and that the absolute energy scale remains model-dependent.

To address the reviewer's concern more directly, we also tested an alternative incident distribution in addition to the Maxwellian assumption, namely a monoenergetic (Dirac-like) distribution, already displayed in Figure 4 from the submitted paper. The result is that the absolute energy scale is shifted, as expected from the different assumed incident spectrum, but the MLT-dependent variability remains qualitatively the same. This indicates that the main dawnside hardening reported here is robust within the framework of the tested transport-model assumptions and is not simply an artifact of choosing a Maxwellian distribution.

Reviewer comment 4:

The AE-based analysis shows that trends are clearer under quiet conditions and more scattered during active periods. This behavior is expected due to increased contamination from injections and discrete aurora during active times. However, the manuscript does not extract new physical insight from this result. In particular, the role of different wave modes (e.g., chorus vs ECH) is not disentangled, and the connection to magnetospheric drivers remains indirect. Using AE alone as a proxy may be insufficient to support the proposed interpretation.

Response:

We thank the reviewer for this important comment. We agree that AE alone is not sufficient to disentangle the relative roles of chorus, ECH, or other scattering processes, and we did not intend to imply that the present dataset allows such a mechanistic separation. We will revise the manuscript to make this limitation more explicit and to adopt a more cautious interpretation of the AE-based analysis.

Our intention in using AE was not to identify the dominant wave mode, but rather to evaluate whether the spectroscopically inferred MLT trend depends on the broader geomagnetic context. In this sense, AE is used here as a practical large-scale proxy of auroral-zone activity and magnetospheric preconditioning, not as a direct diagnostic of the physical scattering mechanism. We agree that the submitted version may have gone too far in a few places by suggesting a more specific physical interpretation than the data can robustly support, and this will be softened in the revised manuscript.

To address related concerns raised by Referee 1, we propose the following revised text in italic for Section 4.2:

4.2 AE over preceding time intervals

To test whether the observed MLT-dependent hardening depends only on the instantaneous geomagnetic context, we also examined the AE index over preceding intervals of 2 h and 6 h relative to each spectrograph measurement. This extends the type of analysis proposed by Hosokawa and Ogawa (2015) to a continuous spectroscopic dataset, while keeping AE as a proxy of geomagnetic context rather than as a diagnostic of the underlying scattering mechanism. The main result is that the post-midnight hardening remains visible even when the activity level is defined from AE over the preceding hours rather than from AE at the exact observation time.

For a 2-h preceding interval, the number of events classified as quiet (AE $<$ 100 nT) drops from 8167 to 2596. This indicates that many spectra recorded under low instantaneous AE are in fact preceded within the previous two hours by enhanced geomagnetic activity. As a consequence, the Bayesian posterior probability for a positive slope in the AE $<$ 100 nT class decreases from 0.996 (lag = 0 h) to about 0.75, while the class $100 \leq AE < 300$ *nT shows a strengthened slope with* $P(\beta_2 > 0) \approx 0.98$*. This redistribution does not contradict the main trend; rather, it shows that many intervals with low instantaneous AE do not correspond to genuinely quiet conditions, but to a magnetospheric context still influenced by earlier activity. Intervals that remain in the AE $<$ 100 nT class even with a 2-h preceding interval therefore represent weakly preconditioned conditions, for which the positive energy–MLT trend is less clearly expressed.*

For a 6-h preceding interval, consistent with the timescale used by Hosokawa and Ogawa (2015) to characterize preceding substorm activity, the redistribution of events among AE classes becomes even more pronounced. In particular, the number of events classified as AE $<$ 100 nT is strongly reduced, indicating that most intervals with low instantaneous AE are preceded by enhanced geomagnetic activity within the previous several hours. At this timescale, the AE-based classification becomes less discriminant, especially for the quiet class. Nevertheless, the post-midnight hardening remains visible. The trend is strongest for intermediate AE levels, with $P(\beta_2 > 0) \approx 0.997$*, and more dispersed for highly active periods (*$P(\beta_2 > 0) \approx 0.83$*).*

Overall, these results show that the MLT-dependent hardening is not captured by instantaneous AE alone, and that its statistical expression depends on the broader geomagnetic context over timescales of several hours. In this sense, the AE analysis is used here to test the robustness of the spectroscopic MLT trend under different activity conditions, rather than to distinguish between specific precipitation mechanisms.

References

Adachi, K., Nozawa, S., Ogawa, Y., Brekke, A., Hall, C., and Fujii, R.: Evaluation of a method to derive ionospheric conductivities using two auroral emissions (428 and 630 nm) measured with a photometer at Tromsø (69.6°N), Earth, Planets and Space, 69, 90, https://doi.org/10.1186/s40623-017-0677-4, 2017.

Fukizawa, M., Sakanoi, T., Miyoshi, Y., Hosokawa, K., Shiokawa, K., Katoh, Y., Kazama, Y., Kumamoto, A., Tsuchiya, F., Miyashita, Y., Tanaka, Y. M., Kasahara, Y., Ozaki, M., Matsuoka, A., Matsuda, S., Hikishima, M., Oyama, S., Ogawa, Y., Kurita, S., and Fujii, R.: Electrostatic Electron Cyclotron Harmonic Waves as a Candidate to Cause Pulsating Auroras, Geophysical Research Letters, 45, 12,661–12,668, https://doi.org/https://doi.org/10.1029/2018GL080145, 2018.

Fukizawa, M., Sakanoi, T., Miyoshi, Y., Kazama, Y., Katoh, Y., Kasahara, Y., Matsuda, S., Matsuoka, A., Kurita, S., Shoji, M., Teramoto, M., Imajo, S., Sinohara, I., Wang, S.-Y., Tam, S. W.-Y., Chang, T.-F., Wang, B.-J., and Jun, C.-W.: Pitch-Angle Scattering of Inner Magnetospheric Electrons Caused by ECH Waves Obtained With the Arase Satellite, Geophysical Research Letters, 47, e2020GL089 926, https://doi.org/https://doi.org/10.1029/2020GL089926, e2020GL089926 2020GL089926, 2020.

Grono, E. and Donovan, E.: Differentiating diffuse auroras based on phenomenology, Annales Geophysicae, 36, 891–898, https://doi.org/10.5194/angeo-36-891-2018, 2018.

Hosokawa, K. and Ogawa, Y.: Ionospheric variation during pulsating aurora, Journal of Geophysical Research (Space Physics), 120, 5943–5957, https://doi.org/10.1002/2015JA021401, 2015.

Kawamura, S., Nishiyama, T., Kataoka, R., Ogawa, Y., Sato, Y., Tsuda, T., Kurita, S., and Saito, Y.: Energy characteristics of pulsating auroras derived from all-sky camera observations, Earth, Planets and Space, 72, 168, https://doi.org/10.1186/s40623-020-01229-8, 2020.

Liang, J., Gillies, D. M., Spanswick, E., and Donovan, E. F.: Converting TREx-RGB green-channel data to 557.7 nm auroral intensity: Methodology and initial results, Earth and Planetary Physics, 8, 258–274, https://doi.org/10.26464/epp2023063, 2024.

Partamies, N., Whiter, D., Kadokura, A., Kauristie, K., Nesse Tyssøy, H., Massetti, S., Stauning, P., and Raita, T.: Occurrence and average behavior of pulsating aurora, Journal of Geophysical Research: Space Physics, 122, 5606–5618, https://doi.org/https://doi.org/10.1002/2017JA024039, 2017.

Robert, E., Barthelemy, M., Cessateur, G., Woelfflé, A., Lamy, H., Bouriat, S., Gullikstad Johnsen, M., Brändström, U., and Biree, L.: Reconstruction of electron precipitation spectra at the top of the upper atmosphere using 427.8 nm auroral images, Journal of Space Weather and Space Climate, 13, 30, https://doi.org/10.1051/swsc/2023028, 2023.

Whiter, D., Rae, J., Samaddar, S., Krcelic, P., Lanchester, B., Carlyle, I., Brindley, N., Thillaimaharajan, K., and Fear, R.: Embla: A new optical instrument to measure auroral precipitation, neutral temperature and electric fields at high resolution, in: EGU General Assembly 2024, pp. EGU24–19 428, https://doi.org/10.5194/egusphere-egu24-19428, 2024.